HPC in Manufacturing using Digital Twins with NVIDIA

How Lenovo® and NVIDIA® are accelerating CAE and AI-enabled Product Development Workflows

2023

Business Whitepaper

Connection
we solve IT

Introduction

Manufacturing is poised to make a giant leap into the future with new technologies and a little boost for some traditional standards. NVIDIA OmniverseTM Enterprise enables manufacturers to take CAE to the next level and embrace digital twins by integrating CAE with AI/ML and other digital technologies. With digital twins, manufacturers can foster organization-wide collaboration and replace many expensive, cumbersome, manual processes with digital ones to optimize the entire product lifecycle.

The Lenovo EveryScale OVX Solution for NVIDIA Omniverse, built upon the Lenovo Scalable Infrastructure (LeSI) framework and ThinkSystemTM SR670-V2 server, is a unique flagship offering for large-scale digital twins resulting from a deep partnership with NVIDIA. As a result, it is the ideal platform to accelerate CAE and the entire product lifecycle. These solutions include:

- NVIDIA graphics processing units (GPUs) to significantly accelerate computer-aided engineering (CAE) solvers
- Lenovo's extensive, expertly engineered portfolio of HPC systems with NVIDIA GPUs helps manufacturers of all sizes take their innovation journey to the next level
- Lenovo and NVIDIA automated and accelerated solutions for CAE post-processing. These include high-performance workstations, remote visualization solutions, and NVIDIA Omniverse.

Technology and Digital Twins vital for manufacturing

To drive innovation and compete aggressively, manufacturers integrate several traditional and emerging technologies (Figure 1) across the product lifecycle into a digital workflow. Key technologies include Computer-Aided Design/Engineering/Manufacturing (CAD/CAE/CAM), Augmented/Virtual Reality (AR/VR), 3D Printing, Robotics, the Internet of Things (IoT), Artificial Intelligence (AI)/Machine Learning (ML), and Digital Twins.

A digital twin is a digital replica, or "twin," of a physical object or business process. Digital twins can be digital replicas of real 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, and gain a competitive advantage. Consequently, the worldwide market for digital twin platforms is expected to grow from $3.2 billion in 2020 to $184.5 billion by 2030.[1]

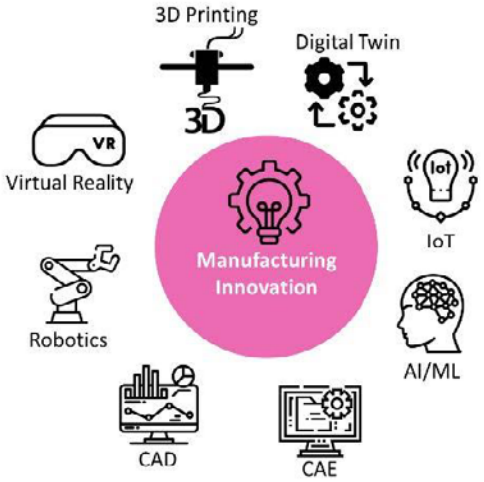

Figure 1: Key technologies driving manufacturing innovation

[1] Technology Trends 2022 | Tech Vision | Accenture

The end to end digital workflow in manufacturing

The end-to-end “digital workflow” (Figure 2) integrates the entire product lifecycle: virtual product development, production, and customer operations. In addition, the potential for Knowledge Feedback through real-time data and analytics opens new avenues for monitoring and improving processes throughout the product life cycle.

In the virtual product development phase, manufacturers typically use physics-based CAD/CAE applications from many independent software vendors (ISVs) such as Ansys®, Siemens®, Dassault Systèmes®, and Altair®. Then, in the production phase, they use various types of real-time operational data with CAM, analytics, and ML to improve supply chain management (SCM), product quality and optimize all production assets. Finally, in the customer operations stage, manufacturers can leverage real-time data with AI/ML to monitor the actual use of their products and provide highly tailored services for predictive maintenance, pricing optimization, and more.

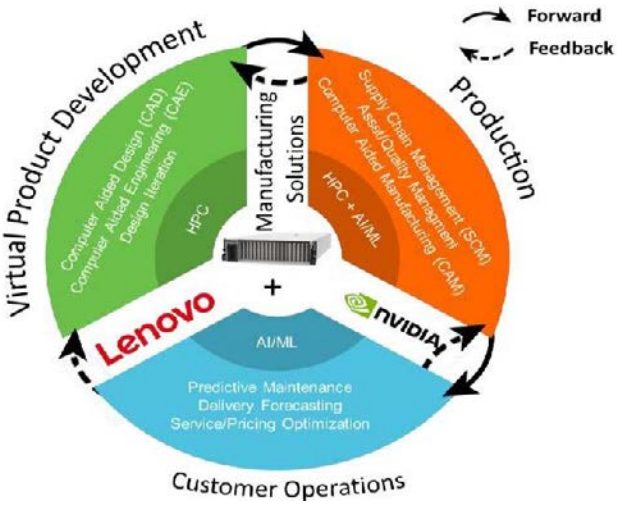

Figure 2: The end-to-end digital workflow for the product lifecycle

Using and integrating CAD, CAE, CAM, and AI/ML functions and applications detailed in Figure 2 can improve productivity and time-to-market and deliver the highest quality products at the lowest costs. Consequently, CAE and AI/ML applications continue to grow, and manufacturers want to streamline and significantly accelerate these workloads.

Accelerating CAE and AI with Lenovo and NVIDIA

CAE has been used in product design and development for decades to shorten development time, optimize product designs, evaluate new design ideas, and minimize costly design changes. However, CAE is now experiencing additional growth because of the increasing use of thermal, structural, electronics, computational fluid dynamics (CFD), and other interdisciplinary analyses in designing and developing more complex products.

More recently, ML methods based on neural networks (known as Deep Learning, or DL) on clusters with graphics processing units (GPUs) are accelerating the CAE process and improving its accuracy by iteratively refining the insights gained from multiple CAE simulation runs that are typical in design optimization loops.

This significantly automates the CAE process, makes complex engineering tools easier to use and opens up new manufacturing business models. However, with the continuing explosion of data, the growing use of multi-physics simulations, and the need to better integrate across various manufacturing silos, including suppliers, there is a need to accelerate all computational processes even more.

Manufacturers can perform more realistic, interdisciplinary simulations with accelerated CAE systems to drive more significant innovation and productivity. Together, NVIDIA and Lenovo deliver these accelerated systems.

The NVIDIA A100 Tensor Core GPU delivers unprecedented acceleration—at every scale. The A100 can efficiently scale up with seven isolated GPU instances/partitions. In addition, multi-instance GPU (MIG) provides a unified platform to adjust dynamically to shifting workload demands.

Using the NVIDIA A100 Tensor Core GPU, Lenovo provides highly reliable, scalable, and significantly accelerated CAE and AI/ML solutions (Table 1). This Lenovo portfolio of servers includes ThinkSystem SR650, ThinkSystem SR655, ThinkSystem SR665, ThinkSystem SR860, and ThinkSystem SR670. In addition, these Lenovo systems with NVIDIA A100 GPUs can perform many parallel CAE computations much faster than traditional central processing units (CPUs).

CAE	GPUs	Lenovo ThinkSystem servers for CAE with NVIDIA GPUs
Ansys Fluent Ansys Mechanical Ansys Rocky DEM	 NVIDIA A100	
SIMULIA Abaqus		
Altair DEM Altair NANOFLUIDX		
Siemens Simcenter STAR-CCM+		

Table 1: Key CAE Applications accelerated by Lenovo ThinkSystem servers enhanced with NVIDIA GPU

Table 2 shows the significant performance improvement of prominent ISV CAE solvers with GPUs over CPUs based on NVIDIA testing in conjunction with the respective ISV partner. All multi-GPU cases use PCIe. These benchmarks are described in greater detail in the appendix.

CAE Application	Speedup range depending on the problem
Ansys Fluent	5.23 on 1 GPU to 32.95 on 8 GPUs over a CPU
Ansys Mechanical	1.3 to 6.8 on 1 GPU over a CPU
Ansys Rocky DEM	40 on 1 GPU to 90 on 4 GPUs over a CPU
SIMULIA Abaqus	9.5 on 1 GPU over a CPU
Altair DEM	1.68 on 2 GPUs to 5.28 on 8 GPUs over 1 GPU
Altair NANOFLUIDX	1.9 on 2 GPUs to 7.7 on 8 GPUs over 1 GPU
Siemens Simcenter STAR-CCM+	6.18 on 2 GPUs to 20.78 on 8 GPUs over 1 CPU

Table 2: Significant performance improvement of CAE solvers on GPUs

While CAE solver acceleration is critical, the infrastructure must also support CAE post-processing needs to improve the engineer's productivity and insights. Lenovo and NVIDIA also provide this infrastructure.

Lenovo + NVIDIA post-processing solutions for design and simulation

CAE post-processors import solver output files and produce comprehensive reports with impactful visualization and charts. It is typically a laborious and time-consuming step and worsens as the results data get larger. Lenovo and NVIDIA provide high-performance infrastructure solutions to automate and accelerate this critical step to enhance the engineer's productivity. In addition, Lenovo offers a broad portfolio from high-powered workstations to scalable, GPU-rich servers, which, combined with NVIDIA GPUs, offer wide-ranging solutions for CAE post-processing. Key members of the portfolio include:

- Lenovo Workstations: ThinkStation P Series workstations deliver powerful performance: from the ThinkStation P620 with its AMD RyzenTM Threadripper PRO to the latest generation of Intel® Xeon® processors combined with up to 2x NVIDIA RTX™ A6000 graphics cards in dual- and single-processor systems. These are ISV-certified, energy-efficient, and highly versatile.
- Lenovo remote visualization: Provides reliable and secure access to graphics-intensive applications anytime, anywhere, from less expensive 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 offers this solution by partnering with Mechdyne®.

CAE solver acceleration and post-processing are crucial. However, these activities are sometimes ad-hoc, uncoordinated, and distributed across the enterprise from other product development activities. Digital twins fix many of these issues.

Leadership solutions for large-scale Digital Twins

The potential for new insights and greater efficiencies is enormous if it is possible to connect the digital processes seamlessly throughout the product lifecycle to mirror the production process. This mirrored connected process is a Digital Twin. Figure 3 depicts this for vehicle manufacturing.

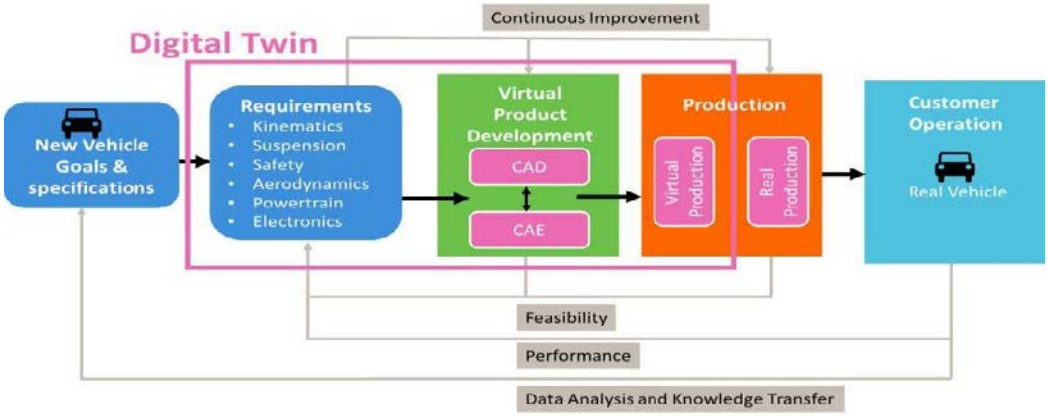

Figure 3: High-level end-to-end development and manufacturing process and the role of the Digital Twin

In partnership with NVIDIA, Lenovo delivers top industry 3D simulation and design collaboration capabilities with the Lenovo EveryScale OVX solution, a purpose-built NVIDIA Omniverse computing system designed to run complex physics and AI-based simulations and industrial-scale digital twins. NVIDIA OVX is a reference architecture for a computing system designed to power the creation and operation of massively complex models and full-fidelity, physically accurate simulation environments in real-time, leveraging the NVIDIA Omniverse platform.

Omniverse Enterprise enables enterprise teams to achieve seamless 3D design collaboration or design and operate real-time physically accurate digital twins. Managing these complex, large-scale digital twins requires extremely high compute, high throughput, low latency, and precision timing.

Lenovo AI infrastructure is purpose-built to meet the demands of large-scale digital twins, optimized to run on Lenovo NVIDIA-certified systems. Lenovo is the only infrastructure solutions provider to deliver the complete portfolio of solutions, from workstations to the edge to the data center, necessary to remotely interface with highly complex digital twins and power 3D collaboration in real-time. The infrastructure delivers:

- Optimized infrastructure for digital twin deployments: NVIDIA OVX on Lenovo ThinkSystem SR670 V2 provides the features and performance required to build large-scale digital twins, physically accurate, AI-enabled replicas of physical environments and processes that synchronize with real-time data. Combination of high-performance GPU accelerated graphics and AI capabilities, high-speed, secure networking, and enterprise-grade management in Lenovo's leading data center servers.
- Scalable performance: The OVX computing system's architecture is designed for multi-node scalability to support the most demanding Omniverse workloads. Lenovo SR670 V2 server allows multiple CPU/GPU configurations to meet the requirements of Omniverse/OVX. In addition, it allows for scaling from a single POD consisting of 8 Lenovo SR670 V2 servers to much larger configurations. Management systems, switches, and storage, up to a superpod for massive digital twin simulations.
- Remote work capabilities: Lenovo and NVIDIA partnership solutions deliver real-time remote work capabilities enabled by NVIDIA Omniverse Enterprise running on the latest NVIDIA GPU-equipped Lenovo platforms. Omniverse Enterprise provides a flexible, customizable 3D design collaboration platform to build digital twins.

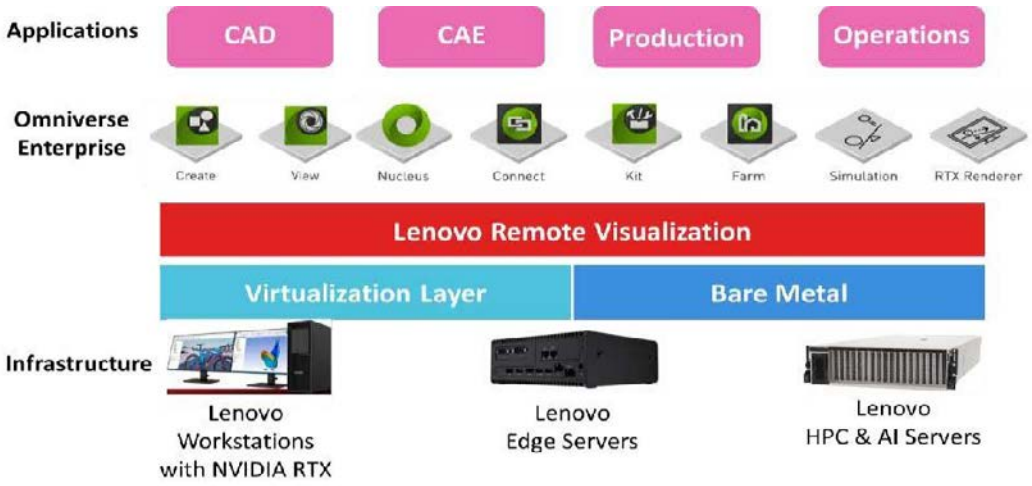

Figure 4: High-level end-to-end development and manufacturing process and the role of the Digital Twin

The details of the solution by each layer (Figure 4) are as follows:

Applications: Product Design, engineering, manufacturing, scientific computing, robotics, and industrial digital twins applications are available for Lenovo and NVIDIA Omniverse environments.

Omniverse Enterprise: NVIDIA Omniverse Enterprise is a scalable, end-to-end platform enabling enterprises to build and run metaverse applications. Metaverse comprises multiple technologies and is the next level of interaction in both virtual and physical worlds that provides innovative new opportunities and business models by allowing businesses to extend digital business to be persistent, decentralized, collaborative, and interoperable. For example, by using digital twin technologies, automotive dealerships can reduce the inventory they have on hand by using spatial computing to change the interior/exterior attributes digitally in real-time. Based on Universal Scene Description (USD), Omniverse Enterprise allows teams to connect and customize complex 3D pipelines and operate large-scale, physically accurate virtual worlds, complex simulations, and industrial digital twins. In addition, it unites design teams, assets, and software tools, enabling diverse workgroups to iterate collaboratively and interactively review photoreal, physically accurate 3D projects and scenes. The critical platform components are:

- Create: For technical artists, designers, and engineers, Omniverse Create accelerates advanced scene composition and allows users to assemble, light, simulate, and render scenes in Pixar USD in real time.
- View: Powers seamless collaborative design and immersive visualization of design and simulation projects for reviewers, clients, and project managers.
- Nucleus: Lets users store, share, and collaborate on project data, providing the unique ability to collaborate live across multiple applications.
- Connect: Allows users to connect leading industry tools to the Omniverse platform for live-sync workflows and save USD and MDL content.
- Kit: A powerful toolkit for developers to build custom tools and plugins quickly and easily in Python or C++ or custom UIs to accelerate design workflows.
- Farm: Developers can quickly build tools to automate repetitive tasks, from rendering workflows to creating turntables, generating thumbnails, synthetic data generation, and more.
- Simulation: Leverage advanced NVIDIA physics technologies for physically accurate simulation.
- RTX renderer: Visualization of scenes in full fidelity with the advanced, multi-GPU RTX Renderer that supports real-time ray tracing, interactive path tracing, and accurate NVIDIA Iray rendering.

Virtualization and bare metal: Traditional CAE solvers use multiple bare-metal physical servers to get the best performance, whereas Remote Visualization and CAD sessions benefit from virtualization.

Infrastructure: Lenovo offers a broad portfolio of servers and workstations, which, combined with NVIDIA GPUs, provides wide-ranging solutions to the manufacturing industry. In addition to the workstations and the servers with NVIDIA A100 Tensor Core GPUs discussed earlier, this portfolio consists of:

- Lenovo Edge Servers: Deliver purpose-built and secure platforms suitable for compute-intensive and latency-sensitive applications deployed outside traditional data centers. They range from purpose-built, compact, and secure ThinkEdge SE-350 to the latest AI server for the edge – ThinkEdge SE-450.
- Lenovo ThinkSystem SR670-V2: It is a versatile GPU-rich 3U rack server that supports eight double-wide GPUs, including the NVIDIA A100 and A40 Tensor Core GPUs or the

NVIDIA HGX A100 4-GPU. It offers NVLink and Lenovo Neptune hybrid liquid-to-air cooling. The server has the new third-generation Intel Xeon Scalable processor and the latest Intel OptaneTM Persistent Memory 200 Series.

Some models of ThinkSystem SR670-V2 include a model with the Lenovo Neptune™ hybrid cooling module, which quickly dissipates heat in a closed-loop liquid-to-air heat exchanger (L2A), delivering the benefits of liquid cooling without adding plumbing.

The Lenovo and NVIDIA advantage for manufacturing

As global technology leaders, Lenovo and NVIDIA spur innovation better by scaling to more significant numbers of parallel tasks than would be practical in CPU-only systems.

Lenovo and NVIDIA advantages for manufacturing:

- Unique OVX solution for digital twins: As a premier OVX partner, the Lenovo EveryScale OVX solution, co-designed with NVIDIA, provides an end-to-end collaboration and true-to-reality simulation platform that transforms complex digital twin workflows at any scale.
- Scalable performance: It features multi-node scalability for the most demanding CAE applications from a vibrant and growing ecosystem of ISVs. Clients can optimize multiple CPU/GPU configurations with remote work capabilities for various digital manufacturing workflows.
- Enterprise-level support: These systems are tested, validated and optimized for performance, manageability, security, and scalability. Lenovo, or a certified business partner, provides on-site installation, start-up, integration, and proactive monitoring and remediation of any CAE operational issues.
- A complete portfolio of solutions: This gives customers an easy path from workstations to the most scalable ThinkSystem Servers for HPC. 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 cloud.

- Solid roadmap with continuing innovation: The latest NVIDIA H100 Tensor Core GPU is designed to deliver an order-of-magnitude performance leap for large-scale AI and HPC over the prior-generation NVIDIA A100 Tensor Core GPU. This GPU is also very energy efficient. The Lenovo ThinkSystem SR670 V2 with NVIDIA H100 GPUs tops the Green500 list of high-performance systems.

Finally, as manufacturers of all sizes progress on their digital transformations, they need reliable partners like Lenovo and NVIDIA with deep CAD, CAE, CAM, and AI/ML expertise throughout their journey.

Appendix: Key GPU-accelerated CAE applications

The following sub-sections describe comparative performance improvement of critical CAE ISV applications with GPUs over CPUs based on NVIDIA testing in conjunction with the respective ISV partner. In addition, all multi-GPU cases use PCIe.

Ansys® Fluent™: A general-purpose computational fluid dynamics (CFD) software that models fluid flow, heat and mass transfer, chemical reactions, and more. Fluent offers a modern, user-friendly interface streamlining the CFD process from pre- to post-processing within a single window workflow. In addition, fluent can leverage a multi-GPU solver on one or multiple GPUs for incompressible and compressible flows and steady-state or transient simulations. This new solver has been architected from the ground up to run natively on GPUs and exploit their full potential.

"The combined solution of Ansys, NVIDIA, and Lenovo provides customers with an HPC environment that is optimized for running engineering simulations. NVIDIA GPUs provide significant performance speedups for many Ansys workloads. This helps enable our customers to conduct more predictively accurate, realistic multi-physics simulations during the design process faster than ever before to drive innovative new product designs.", Wim Slagter, Director Strategic Partnerships, Ansys

Figure 5 shows the performance improvement with the NVIDIA GPUs. Compared to a dual CPU system as the base, the performance improvement for one NVIDIA A100 GPU and eight of the same GPU ranges from 5-fold to 32-fold, respectively.

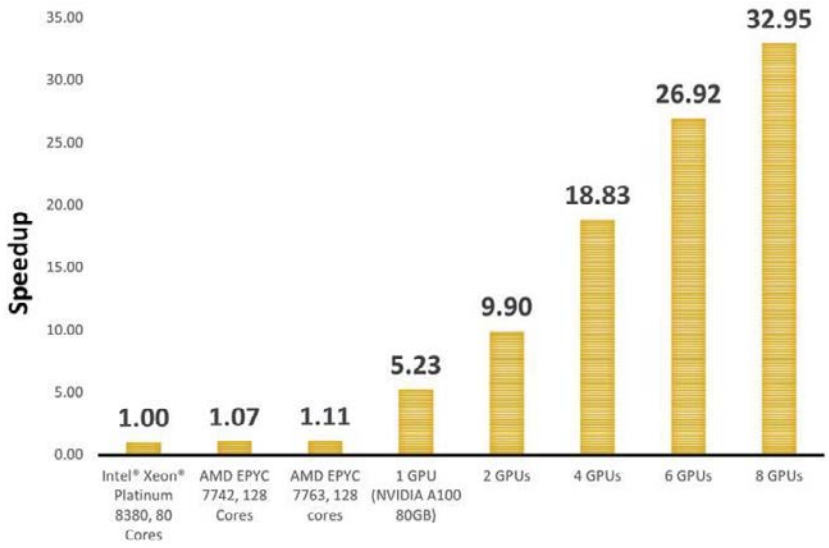

Figure 5: Performance gains with NVIDIA testing for Ansys Fluent solver on DrivAer model with NVIDIA GPUs

Ansys® Mechanical™: Engineers can solve complex structural engineering problems and make better, faster design decisions. With the finite element analysis (FEA) solvers available in the suite (Table 3), users can customize and automate solutions for structural mechanics problems and parameterize them to analyze multiple design scenarios. In addition, NVIDIA GPU acceleration enables more efficient computation and faster job turnaround times, delivering more license utilization for the same investment.

Figure 6 shows the benefits of combining NVIDIA A100 GPUs with conventional CPUs. Depending on the problem (Table 3), the improvement in the total solution time ranges from 1.3 times to 6.8 times.

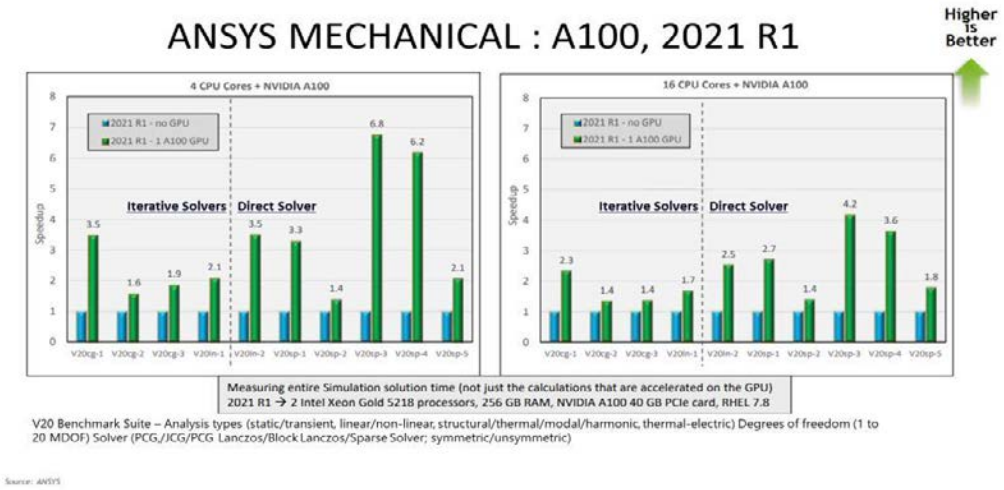

Figure 6: Performance improvement based on NVIDIA testing for Ansys Mechanical with NVIDIA GPUs

Problem Description	Solver	Matrix Type	DOF	Analysis Type
Power Supply Module V20cg-1	JCG solver	Symmetric	6.7 DOF	Thermal analysis, Static & Linear
Tractor Rear Axle V20cg-1	PCG solver	Symmetric	19.2 DOF	Structural analysis, Static & Linear
Engine block V20cg-3	PCG Lanczos Eigen Solver	Symmetric	23.3 DOF	Structural analysis, Static & Linear
Gear Box V20ln-2	Block Lanczos Eigen Solver	Symmetric	13.9 DOf	Structural analysis, Modal & Linear
Radial Impeller V20ln-2	Sparse solver	Symmetric	4.4 DOF	Structural analysis, Cyclic symmetry & Modal
Peltier Cooling block V20sp-1	Sparse solver	Symmetric	1.3 DOF	Thermal electric coupled field analysis, Static & Nonlinear
Semi submersible V20sp-2	Sparse solver	Symmetric	8.3 DOF	Structural analysis, Transient & Nonlinear
Speaker V20sp-3	Sparse solver	Symmetric	2.5 DOF	Structural analysis, Harmonic & Linear
Turbine V20 sp-4	Sparse solver	Symmetric	5.9 DOF	Structural analysis, Static & Nonlinear
BGA V20sp-5	Sparse solver	Symmetric	16.1 DOF	Structural analysis, Transient & Nonlinear

Table 3: Significant performance improvement of CAE solvers on GPUs

Ansys® Rocky DEM™: a high-fidelity particle simulation software that quickly and accurately simulates the flow behavior of particles with complex particle shapes and size distributions to simulate material flows and machining processes realistically. According to Ansys, nearly 70

percent of industrial products experience bulk granular material flows, where different-sized particles with complex shapes interact, potentially impacting a product's efficiency or structural integrity. Rocky DEM is integrated with Ansys Fluent and Ansys Mechanical and simulates fully coupled fluid and structural systems' bulk material movement and reaction to forces such as gravity and adhesion.

Rocky DEM uses multiple GPUs simultaneously to analyze bulk material flow systems, and as Figure 7 shows, the speed-up can range from 40 to 90 times.

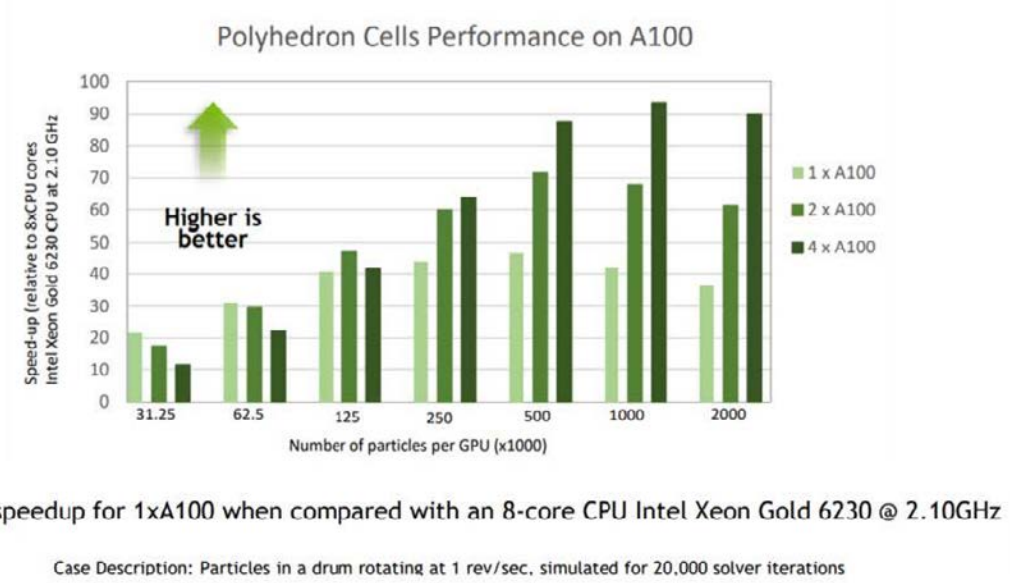

Figure 7: Performance improvement based on NVIDIA testing for Rocky DEM with NVIDIA GPUs

Siemens Simcenter STAR-CCM+™: A Multiphysics CFD software for the simulation of products operating under real-world conditions. It includes CAD, automated meshing, Multiphysics CFD, sophisticated post-processing, and design exploration. This integrated environment allows engineers to efficiently explore the entire design space to make better

design decisions faster. For Simcenter STAR-CCM+ 2022.1 software, eight NVIDIA A100 GPUs provide up to 20 times faster (Figure 8) than conventional CPU servers.

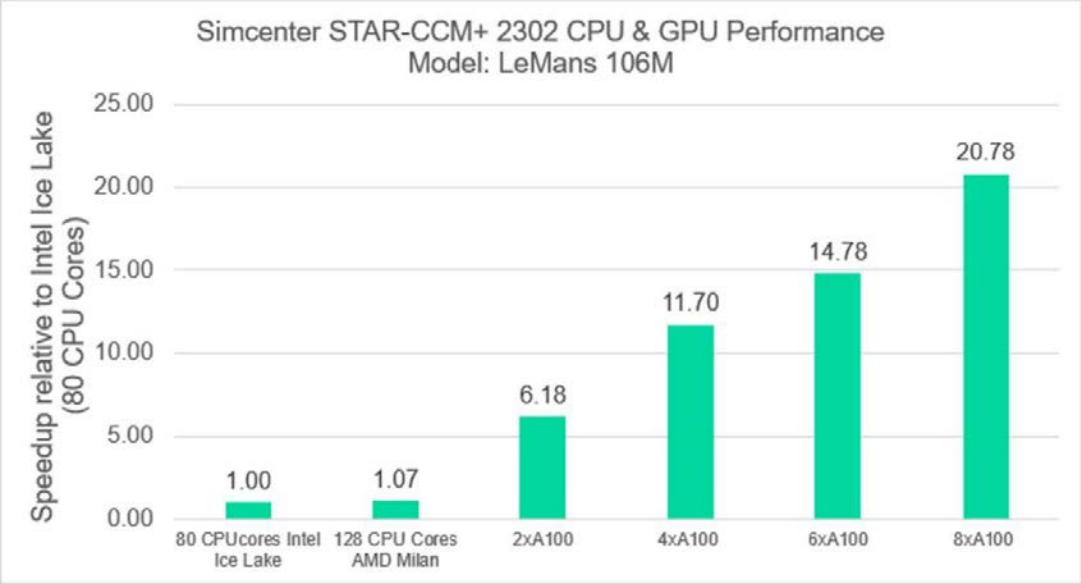

Figure 8: Performance improvement based on NVIDIA testing for Simcenter STAR-CCM+ with NVIDIA GPUs

"The joint solution of Siemens Simcenter STAR-CCM+, NVIDIA, and Lenovo provides customers with an optimized CAE environment NVIDIA GPUs provide substantial performance improvements for many STAR-CCM+ workloads. Our customers can conduct more accurate, realistic engineering simulations faster than ever before to drive more innovative product designs.", Liam Mcmanus, Technical Product Manager Siemens Digital Industries Software

SIMULIA Abaqus: The Abaqus Unified FEA product suite from SIMULIA, a Dassault Systèmes brand, offers robust structural and multiphysics solutions for a wide range of industrial applications. For example, in the automotive industry, engineering teams can analyze full vehicle loads, dynamic vibration, multibody systems, impact/crash, nonlinear static, thermal coupling, and acoustic-structural coupling using a common model and unified solver technology. Abaqus/Standard leverages NVIDIA GPU computing to accelerate the analysis of static and low-speed dynamic events, such as sealing pressure in a gasket joint, steady-state rolling of a tire, or crack propagation in a composite airplane fuselage. It enables users to analyze a model in both time and frequency domains within a single simulation. As a result, industry-leading companiesuse the Abaqus suite to consolidate their processes and tools, reduce costs and inefficiencies, and gain a competitive advantage. Figure 9 shows that the NVIDIA A100 with AMD CPU is 9.5 times faster than the AMD CPU.

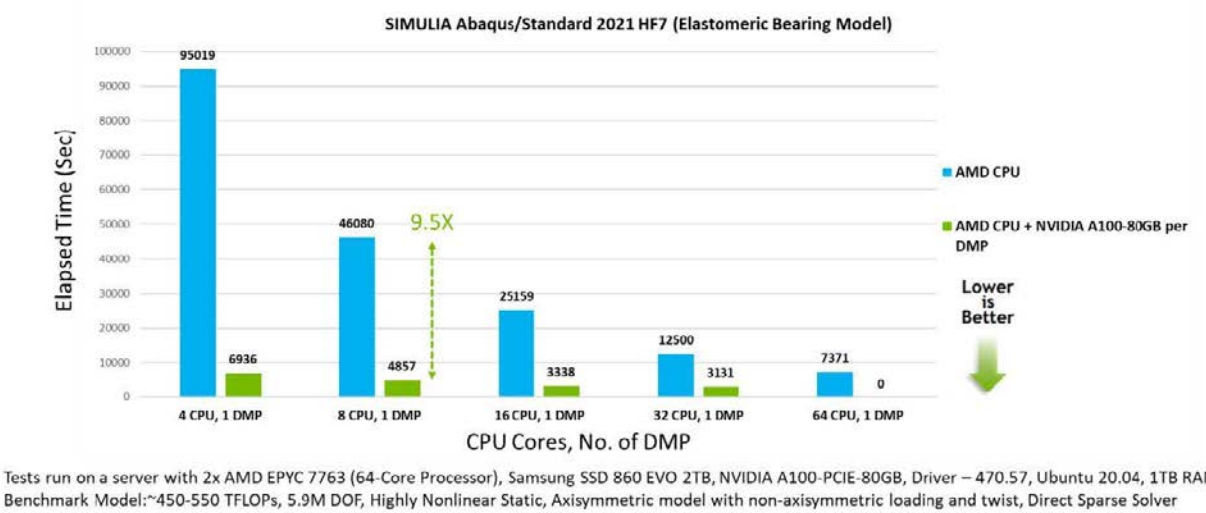

Figure 9: Performance improvement based on NVIDIA testing for SIMULIA Abaqus/Standard with NVIDIA GPUs

"The combined solution of NVIDIA and Lenovo, running SIMULIA realistic simulation applications, provides an optimized HPC environment for design exploration. NVIDIA GPUs provide significant performance speedups for SIMULIA structural and electromagnetic analysis workloads. With this HPC solution, our mutual customers can run high-fidelity Multiphysics simulations faster and accelerate the development of innovative new products", Delphine GENOUVRIER, SIMULIA R&D Roles Portfolio Director, Dassault Systèmes

Altair® nanoFluidX®: A particle-based CFD tool can predict multiphase and single-phase flows in complex geometries with complex motion. It can forecast the oiling in powertrain systems with rotating shafts/gears and analyze forces and torques on individual components of the system. In addition, it can be used for various water management applications, as free surface flows tend to come easily to the particle-based CFD, and GPU performance in such cases stands out.

nanoFluidX was created and optimized for GPU clusters, making it extremely fast. Figure 10 shows the incremental improvement in performance with additional GPUs. Performance improvement ranges from 1.9 to 7.7 times compared to a single GPU solution.

"The combined solution of Altair, NVIDIA, and Lenovo provides customers with an HPC environment that is optimized for running engineering simulations. NVIDIA GPUs provide significant performance speedups for many Altair workloads. This enables our customers to conduct more accurate, large-scale, multi-physics simulations during the design process faster than ever before to drive innovative new product designs.",

David Curry Senior Vice President, CFD and EDEM Products, Altair Engineering

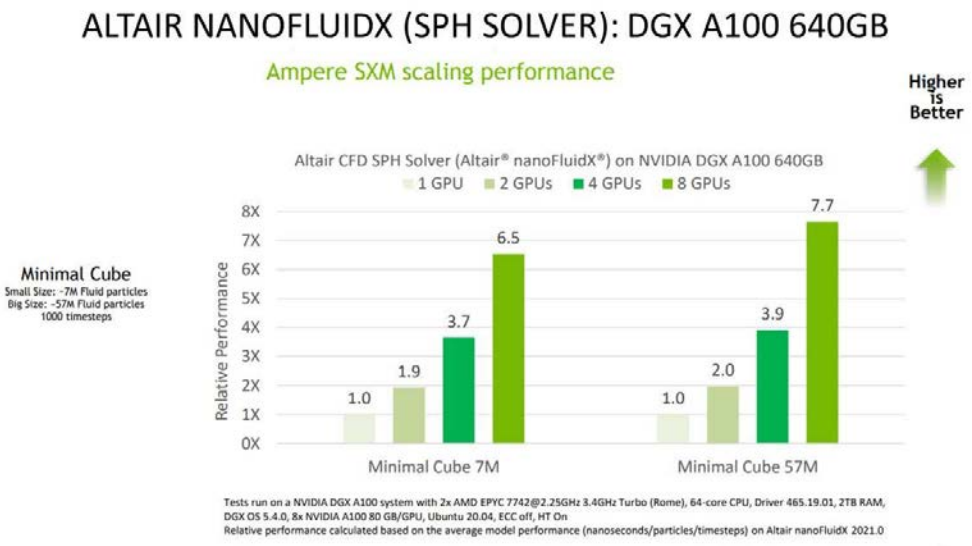

Figure 10: Performance gains with NVIDIA testing for ALTAIR NANOFLUIDX(SPH) solver with NVIDIA GPUs

Altair® EDEM™: A high-performance solution for bulk and granular material simulation. EDEM simulates and analyzes the behavior of coal, mined ores, soils, fibers, grains, tablets, powders, and more. As a result, it is used by engineers from many industries to give crucial insight into how those materials will interact with their equipment during various operations and process conditions. EDEM can be integrated with multiple CFD and multi-body dynamic software and offers 10,000s of ready-to-use material models covering powders to heavy quarry rocks.

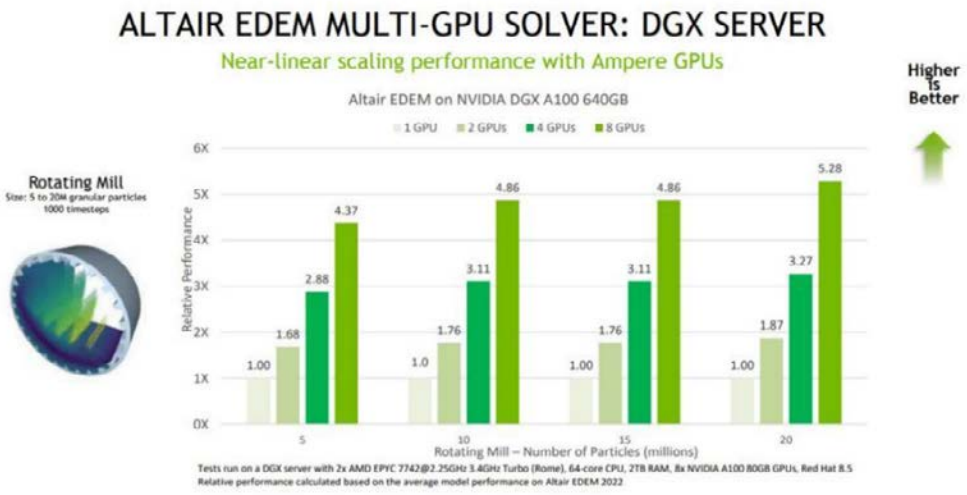

Figure 11: Performance improvement based on NVIDIA testing for ALTAIR EDEM solver with NVIDIA GPUs.

Figure 11 shows the incremental improvement in the performance of EDEM with additional GPUs. Performance improvement ranges from 1.68 to 5.28 times compared to a single GPU solution. 4 GPUs deliver a performance increase of more than 300x over a 32-core CPU[2].

[2] Altair Resource EDEM Benchmark

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to: Lenovo (United States), Inc., 8001 Development Drive, Morrisville, NC 27560, U.S.A. Attention: Lenovo Director of Licensing.

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. This information could include technical inaccuracies or typographical errors. Changes may be made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any performance data contained herein was determined in a controlled environment; therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems, and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Lenovo, the Lenovo logo, System x, and ThinkServer, are trademarks of Lenovo in the United States, other countries, or both.

Intel and Xeon are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows Storage Server 2012, Windows Server 2012, and the Windows Logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

NVIDIA, the NVIDIA logo, NVIDIA OVX, NVIDIA Omniverse, are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries.

Mechdyne® VMware® Citrix® Red Hat® AMD Ryzen™ Intel® Xeon® Intel Optane™ Lenovo Neptune™ Ansys® Fluent™ Ansys Mechanical™ Ansys Rocky DEM™ Siemens Simcenter STAR-CCM+™ Dassault Systèmes SIMULIA® Abaqus® Dassault Systèmes SIMULIA® CST Studio Suite® Dassault Systèmes SIMULIA® PowerFLOW® Altair® nanoFluidX® Altair® EDEM™ Other company, product, or service names may be trademarks or service marks of others.