

ACCELERATING CAE SIMULATIONS WITH ALTAIR ONE® AND MICROSOFT AZURE®

In the race to bring better products to market faster, design and engineering teams rely on multi-physics simulations. Increasingly, these simulations involve state-of-the-art solvers and AI tools that require high-performance GPUs. With Altair One and Microsoft Azure, designers and engineers can run simulations using industry-leading tools and the latest GPU-accelerated cloud instances, boosting performance, reducing cycle times, and delivering outstanding productivity gains.

The Case for Running CAE Simulations in the Cloud

Increasingly, manufacturers compete on the strength of their high-performance computing (HPC) infrastructure. They rely on systems with powerful CPUs and GPUs, large memory capacities, and high-speed networks to support demanding simulation and AI workloads. Deploying and managing HPC infrastructure presents several challenges, including high capital expenditures (CapEx), long procurement cycles, the need for skilled personnel, and physical constraints such as limited floor space, power, and cooling capacity.

Not surprisingly, many organizations are turning to the cloud to help address infrastructure-related challenges. By shifting HPC and AI workloads to the cloud, organizations can avoid the pitfalls of on-premises deployments while realizing multiple benefits:

- On-demand access to the latest CPUs and GPUs
- Faster deployment times for compute and storage
- Seamless scaling to keep pace with changing business needs
- Predictable pricing with flexible, consumption-based cost models
- Avoid tying up scarce capital on expensive infrastructure

However, access to scalable HPC infrastructure is only part of the problem. Manufacturers also need to deploy and manage advanced design and simulation tools that facilitate collaboration between design and engineering teams. Fortunately, Altair One provides an ideal solution.

Using Altair One with Microsoft Azure ND-v5 series VMs for CFD demonstrated an average throughput improvement of ~5 to 6x over a single GPU. With these scalable, cloud-based VMs, engineers can iterate more quickly and explore a wider range of design alternatives, resulting in higher-quality products and a faster time to market.

Contact us at:
altair.com/contact-us

Altair One, the Cloud Innovation Gateway

[Altair One](#) is a revolutionary cloud service for collaborative engineering, data engineering, and the development of analytical applications. Built on a robust HPC backbone in the [Microsoft Azure](#) cloud, and decades of simulation, HPC, and AI expertise, Altair One provides engineers access to the tools, data, and computing resources they need across every stage of the product development lifecycle. Designers and engineers can access the Altair One cloud portal and choose from over 200 Altair and third-party design, simulation, analytics, and AI-based tools. They can download and run applications on local servers or clusters, run design and simulation tools on scalable Azure cloud infrastructure, or seamlessly launch simulations to their preferred public or private clouds.

Altair One, the Cloud Innovation Gateway

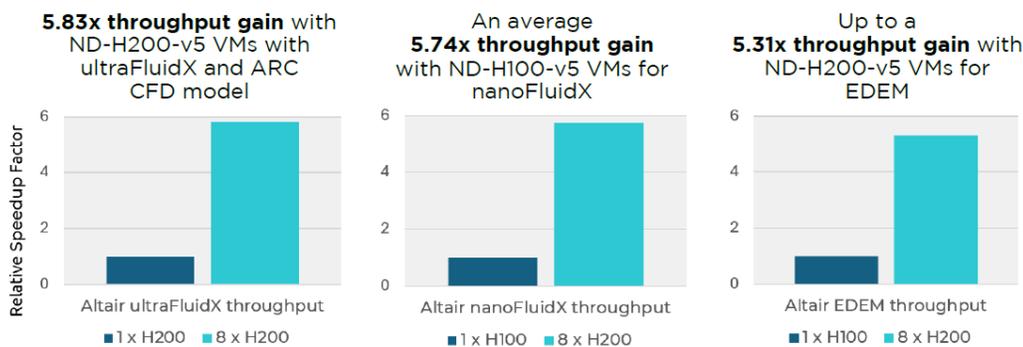
Once considered niche and experimental, GPU-based solvers have been making steady inroads. GPUs are now used to accelerate physics-based simulations in multiple domains, including dynamic explicit analysis, computational fluid dynamics (CFD), and electromagnetics (EM). Traditionally, high-fidelity simulations have run for hours or even days across clusters of high-performance servers. Using the latest NVIDIA® GPUs, engineers can now run these simulations faster and at a lower cost, often on a single server or cloud virtual machine (VM). They can do this while also substantially reducing their energy requirements and associated carbon footprint.

Microsoft offers a broad selection of accelerated VMs, including the [ND-H100-v5](#) and [ND-H200-v5](#) series, designed for large-scale HPC and AI workloads. Ideal for demanding CAE applications, these VMs are based on the NVIDIA® Hopper™ architecture and feature NVIDIA H100 and H200 Tensor Core GPUs, respectively. The ND-H100-v5 series VM (Standard_ND96isr_H100_v5) provides 96 vCPUs, 1,900 GB of system memory, and 8 x NVIDIA H100 Tensor Core GPUs (SXM™ form factor), each with 80 GB of HBM3 memory. Each GPU on the VM is equipped with a dedicated 400 Gb/s NVIDIA Quantum-2 CX7 InfiniBand interface and provides 3.35 TB/s of aggregate memory bandwidth, enabling scalability to thousands of GPUs across multiple nodes.

The ND-H200-v5 series (Standard_ND96isr_H200_v5) offers similar VM-level specs but is powered by 8 x NVIDIA H200 GPUs, each with 141 GB of HBM3e memory and 4.8 TB/s of memory bandwidth for optimal performance.

The Proof is in the Benchmarks

In June 2025, engineers from Altair and Microsoft conducted a comprehensive set of benchmarks to evaluate the performance and scalability of Altair's GPU-accelerated solvers on Microsoft Azure. The benchmarks measured the results of 20 tests with varying numbers of GPUs across 14 models and three GPU-accelerated Altair CAE solvers: [Altair® ultraFluidX®](#) and [Altair® nanoFluidX®](#) for CFD, and [Altair® EDEM™](#), a discrete element modeling (DEM) solver for bulk and granular material simulations. Benchmarks were run on Azure ND-H100-v5 and ND-H200-v5 series VMs, each equipped with eight NVIDIA H100 or H200 Tensor Core SXM GPUs, respectively.



In an ultraFluidX benchmark simulating the external aerodynamics of a vehicle design (Altair ARC SUV model), compute time decreased from 2,744 seconds (45m 44s) on a single H200 GPU to 470 seconds (7m 50s) on eight H200 GPUs. This corresponds to a speedup factor of 5.83x (73% strong scaling efficiency), showcasing the strong scalability of the ultraFluidX solver across multiple GPUs.¹

5.8% ▲

ULTRAFUIDX
THROUGHPUT GAIN
WITH AZURE ND-H200-
V5 SERIES VMs AND
ALTair ARC SUV
MODEL

5.7% ▲

AVERAGE NANOFLUIDX
THROUGHPUT GAIN
WITH ND-H200-V5
SERIES VMs

5.3% ▲

THROUGHPUT GAIN
WITH ND-H200-V5
SERIES VMs RUNNING
EDEM POWDER MIXER
MODEL

Averaged across ten nanoFluidX tests, moving from one H100 to eight H200 GPUs resulted in an average 5.74x speedup, equivalent to a 72% strong scaling efficiency.² This reduction in turnaround time enables engineers to explore more design alternatives within tight project deadlines.

EDEM benchmarks were run on both the Standard_ND96isr_H100_v5 and Standard_ND96isr_H200_v5 VMs for each model with one, two, four, and eight GPUs. Running with all eight GPUs active on the Standard_ND96isr_H100_v5 VM resulted in an average 3.90x speedup relative to a single H100 GPU across the tested models. A slightly higher average uplift of 3.95x speedup was observed when using eight GPUs on the Standard_ND96isr_H200_v5 VM. The best result was achieved with the Powder Mixer model, where the 8 x H200 GPU configuration ran 5.31x faster than a single H100 GPU, reducing the runtime from approximately 45 minutes to 8 minutes, demonstrating a significant productivity improvement.³

Boost Productivity with Azure GPU-Accelerated Instances

Running GPU-accelerated Altair One simulations and AI workloads on Microsoft Azure helps deliver both breakthrough results and outstanding business agility. Together, Altair and Microsoft can help manufacturers realize multiple benefits:

- **Superior time-to-solution** — Across 20 Altair CFD and Altair EDEM test cases, GPU-accelerated Azure VMs slashed runtimes by up to ~5-6x versus a single GPU.
- **Right-sized economics** — Microsoft Azure ND-H100-v5 instances provide outstanding throughput and value for most CFD workloads. The additional memory and bandwidth offered by ND-H200-v5 VMs provide a performance boost and necessary headroom for oversized meshes and AI workloads.
- **Effortless scale, management, and governance** — Altair One and Microsoft Azure abstract the complexity of HPC, enabling engineers to launch jobs with automated infrastructure provisioning in minutes, while IT retains full visibility into spending and resource governance.

Getting Started with Altair One on Microsoft Azure

Whether you're optimizing an EV cooling channel, predicting slurry flow, or training an Altair® PhysicsAI™ model, Altair One and Microsoft Azure give you the speed, scale, and flexibility to iterate faster and design better products—today and into the future.

To learn more about Altair One, visit altairone.com.

To learn more about Microsoft Azure, visit azure.microsoft.com.

RUNNING ALTAIR ONE SIMULATIONS ON AZURE ND-H100-V5 AND ND-H200-V5 SERIES VMS CAN SLASH RUNTIMES FROM HOURS TO MINUTES, POTENTIALLY REPLACING AN ENTIRE CLUSTER.

1. Tests conducted using two ultraFluidX models by Altair and Microsoft in June 2025, comparing the performance of a single H200 GPU to various numbers of GPUs on ND-H100-v5 and ND-H200-v5 series VMs on Azure. Using the proprietary Altair ARC SUV mode, the 8 x H200 configuration delivered 5.83x the throughput of a single H200 GPU.

2. An average of ten tests conducted across four nanoFluidX models by Altair in June 2025, comparing the performance of a single H100 GPU to 8 x H200 GPUs on ND-H100-v5 and ND-H200-v5 series VMs on Azure. The average speed-up across the ten benchmarks was 5.74x.

3. Based on EDEM benchmarks conducted by Altair and Microsoft in June 2025, comparing EDEM scalability across multiple standard models. In the Powder Mixer benchmark, 8 x H200 GPUs on an Azure ND-H200-v5 series VM delivered 5.31x the performance of the baseline single H100 GPU configuration on an Azure ND-H100-v5 series VM.