

Application Notes and Protocols: Leveraging NVIDIA H100 for Computational Fluid Dynamics Research

Author: BenchChem Technical Support Team. **Date:** April 2026

Compound of Interest

Compound Name: H100

Cat. No.: B15585327

[Get Quote](#)

Audience: Researchers, scientists, and drug development professionals.

Introduction

The NVIDIA **H100** Tensor Core GPU, based on the Hopper™ architecture, represents a significant leap in computational power, offering unprecedented acceleration for High-Performance Computing (HPC) and AI workloads.[1] For computational fluid dynamics (CFD), a field critical to disciplines ranging from aerospace engineering to pharmaceutical development, the **H100** provides a transformative tool.[1] Its architectural enhancements, including superior double-precision (FP64) performance, high-bandwidth memory (HBM3), and advanced multi-GPU connectivity, enable researchers to tackle larger, more complex simulations with higher fidelity and in a fraction of the time required by previous hardware generations.[2][3]

These notes provide detailed protocols and performance data for leveraging the **H100** GPU to accelerate CFD research, with a particular focus on applications relevant to the life sciences and drug development.

Section 1: H100 Architecture and its Impact on CFD

The **H100** GPU introduces several key architectural improvements over its predecessors, like the A100, which directly address common bottlenecks in CFD simulations.[2] The shift to a more advanced 4nm process enhances performance per watt, making large-scale simulations more energy-efficient.[2]

Key advantages for CFD include:

- **Enhanced FP64 Performance:** Many traditional CFD solvers require high-precision calculations to ensure robust and accurate solutions.[4] The **H100** offers up to 4x the FP64 performance of the A100, drastically reducing time-to-solution for these essential workloads. [2]
- **High-Bandwidth Memory (HBM3):** CFD simulations, particularly those using methods like the Lattice Boltzmann Method (LBM), are often memory-bandwidth bound.[5][6] The **H100's** HBM3 memory provides up to 3.35 TB/s of bandwidth, enabling faster data access for massive computational grids and reducing solver iteration times.[6][7]
- **Increased Memory Capacity:** With up to 80 GB of memory per GPU, the **H100** can accommodate significantly larger and more detailed meshes in a single node.[3] This minimizes the need for complex domain decomposition across multiple nodes, which can introduce communication overhead.[3]
- **Fourth-Generation NVLink:** For simulations that exceed the memory of a single GPU, the high-speed NVLink interconnect is crucial.[2] It allows multiple **H100s** to function as a unified, powerful accelerator, enabling efficient scaling for massive problems.[2][8]




Diagram 1: H100 Architectural Advantages for CFD

[Click to download full resolution via product page](#)Diagram 1: **H100** Architectural Advantages for CFD

Data Presentation: H100 vs. A100 GPU Specifications

The following table summarizes the key architectural differences between the NVIDIA **H100** and A100 GPUs that are most relevant to CFD workloads.

Feature	NVIDIA A100 (SXM4 80GB)	NVIDIA H100 (SXM5 80GB)	Impact on CFD
Architecture	Ampere (7nm)	Hopper (4nm)	Higher performance per watt.[2]
FP64 Performance	9.7 TFLOPS	67 TFLOPS	Significantly faster double-precision calculations.[9][10]
Memory Type	HBM2e	HBM3	Faster data access for large datasets.[2]
Memory Bandwidth	2 TB/s	3.35 TB/s	Reduces bottlenecks, crucial for bandwidth-bound solvers.[2][7]
Multi-Instance GPU	1st Generation	2nd Generation	More flexible and secure resource provisioning.[1]
NVLink	3rd Generation (600 GB/s)	4th Generation (900 GB/s)	Faster and more efficient multi-GPU scaling.[2]

Section 2: Performance Benchmarks

Quantitative benchmarks demonstrate the substantial performance gains achievable with the **H100** across various industry-standard CFD software packages.

Data Presentation: Ansys Fluent Benchmarks

Ansys Fluent, a widely used CFD solver, shows remarkable speedups on **H100** GPUs. In some cases, a single **H100** GPU can deliver performance equivalent to over 400 CPU cores.[4][11] The metric "MIUPS" (Millions of Iterations and Updates per Second) is used to provide a relative performance measure independent of grid size.[8]

Benchmark Case	# of H100 GPUs	MIUPS	Speedup (vs. 1 GPU)	Parallel Efficiency
car_2m	1	101.38	1.0x	100%
2	119.60	1.18x	59%	
4	154.74	1.53x	38%	
8	184.25	1.82x	23%	
combustor_24m[12]	1	102.32	1.0x	100%
2	196.53	1.92x	96%	
4	385.08	3.76x	94%	
8	743.08	7.26x	91%	

Note: The drop in efficiency for the car_2m case on 8 GPUs suggests the model is too small to scale effectively across that many processors.[\[12\]](#)

Data Presentation: Altair ultraFluidX® and FluidX3D Benchmarks

Altair's GPU-native CFD solver, ultraFluidX, also demonstrates strong performance and scaling on the **H100**.

Benchmark	GPU Configuration	Throughput (vs. A100)	Multi-GPU Scaling Efficiency
Altair Roadster [9]	1x H100 PCIe vs. 1x A100 SXM	~1.4x Higher	N/A
2x H100 PCIe vs. 1x H100 PCIe	N/A	~93% - 97%	

For the bandwidth-bound Lattice Boltzmann solver FluidX3D, the **H100** 80GB PCIe card leads in performance, measured in Mega Lattice Updates per second (MLUPS).[\[5\]](#)

Precision Mode	Performance (MLUPS)	Power Efficiency (MLUP/W)
FP32	~15,500	51.3
FP16S (Scaled FP16)	~22,500	74.5
FP16C (Custom FP16)	~16,300	53.9

Data extracted from performance charts in ServeTheHome's hands-on review.[\[5\]](#)

Section 3: Protocols for Setting Up and Optimizing CFD Workloads

Maximizing the performance of the **H100** requires proper configuration of both the hardware environment and the CFD software.




Diagram 2: General H100-Accelerated CFD Workflow

[Click to download full resolution via product page](#)

Diagram 2: General **H100**-Accelerated CFD Workflow

Protocol 3.1: General Environment Setup

- Hardware Installation:
 - Install NVIDIA **H100** GPUs in servers with PCIe Gen5 slots to ensure maximum host-to-device bandwidth.[\[13\]](#)
 - For multi-GPU setups, utilize motherboards and servers that support NVLink bridges or NVSwitch technology to enable high-speed direct communication between GPUs.[\[2\]](#)
- Software Installation:
 - Install the latest NVIDIA drivers for the **H100** GPUs.
 - Install the NVIDIA CUDA Toolkit. This provides the necessary libraries and compilers (e.g., NVCC) for GPU-accelerated applications.
 - Install the Message Passing Interface (MPI) library (e.g., Open MPI, Intel MPI) required for multi-node and multi-GPU parallel processing.[\[13\]](#)

Protocol 3.2: Optimizing Ansys Fluent Simulations

- Solver Selection: Launch Ansys Fluent and select the appropriate GPU-accelerated solver. Most native Fluent solvers have been optimized for NVIDIA GPUs.
- GPU Activation: In the Fluent launcher or via command-line flags, specify the number of **H100** GPUs to be used for the simulation.
- Precision Mode: While many CFD problems require double precision ("3ddp"), the Fluent GPU solver is robust in single precision ("3d") and may offer faster convergence for certain classes of problems.[\[14\]](#) Test both modes to find the optimal balance of speed and accuracy for your specific case.
- Multi-GPU Execution: For large models that require more than 80 GB of memory, Fluent will automatically distribute the mesh across the specified number of GPUs. The use of NVLink is critical for achieving high parallel efficiency in these cases.[\[8\]](#)

- Monitoring: Use tools like nvidia-smi to monitor GPU utilization, memory usage, and power draw during the simulation to ensure the hardware is being used effectively.

Protocol 3.3: Optimizing OpenFOAM Simulations

- Obtain a GPU-Enabled Version: Use a version of OpenFOAM that has been compiled with CUDA support or use a specialized fork designed for GPUs.[\[13\]](#)
- Solver and Library Configuration:
 - Select solvers that are optimized for GPU execution (e.g., pimpleFoam, icoFoam with CUDA support).[\[13\]](#)
 - Configure the simulation to use NVIDIA's linear algebra libraries like cuBLAS and cuSOLVER for matrix operations to leverage the Tensor Cores.[\[13\]](#)
- Mesh Partitioning:
 - Before running the simulation, decompose the mesh using the decomposePar utility.[\[13\]](#)
 - Choose a decomposition method (e.g., scotch, metis) that creates balanced partitions to ensure an even workload across all participating GPUs.[\[13\]](#)
- Execution Command:
 - Combine MPI with CUDA for parallel execution. Launch the simulation using mpirun, specifying the total number of processes (one per GPU).
 - Example: `mpirun -np 4 simpleFoam -parallel` for a 4-GPU simulation.
- Parameter Tuning:
 - Adjust the deltaT (time step) and maxCo (Courant number) to maintain numerical stability while maximizing the computational work done in each step, thereby keeping the GPU utilization high.[\[13\]](#)
 - Optimize memory usage to ensure the problem fits within the GPU's 80GB of HBM3 memory, minimizing costly data transfers to and from the CPU's system RAM.[\[13\]](#)

Section 4: Application in Drug Discovery and Development

The acceleration provided by **H100** GPUs is particularly impactful in the pharmaceutical industry, where CFD is a vital tool for process development and drug delivery system design. [\[15\]](#)[\[16\]](#)

- **Process Understanding and Scale-Up:** CFD can model complex processes like the hydrodynamics and mass transfer within large-scale pharmaceutical hydrogenation reactors. [\[15\]](#) **H100**-accelerated simulations allow for rapid virtual testing of different impeller speeds, solvent volumes, or reactor geometries, reducing the need for costly and time-consuming physical experiments.[\[15\]](#)
- **Drug Delivery Device Design:** The development of pharmaceutical aerosol devices, such as dry powder inhalers (DPIs) and pressurized metered-dose inhalers (pMDIs), relies on understanding the complex airflow and particle dynamics within the device.[\[16\]](#) High-fidelity CFD simulations on **H100s** can predict aerosol behavior and lung deposition, facilitating the optimization of device design for better therapeutic efficacy.[\[16\]](#)
- **Synergy with AI and Molecular Dynamics:** The drug discovery pipeline heavily relies on other computationally intensive tasks like molecular docking and molecular dynamics simulations, which are also significantly accelerated by GPUs.[\[17\]](#) The **H100** provides a unified platform that can accelerate not only the fluid dynamics aspects of manufacturing and delivery but also the foundational AI-driven and simulation-based discovery of new drug candidates.[\[17\]](#) [\[18\]](#)




Diagram 3: H100-Accelerated CFD in the Drug Development Pipeline

[Click to download full resolution via product page](#)

Diagram 3: **H100**-Accelerated CFD in the Drug Development Pipeline

Need Custom Synthesis?

BenchChem offers custom synthesis for rare earth carbides and specific isotopic labeling.

Email: info@benchchem.com or [Request Quote Online](#).

References

- [1. clustervision.com](https://clustervision.com) [clustervision.com]
- [2. How do the performance and efficiency of A100 and H100 GPUs compare in large-scale fluid dynamics simulations? - Massed Compute](#) [massedcompute.com]
- [3. How does the increased memory capacity of H100 GPUs impact performance in CFD simulations? - Massed Compute](#) [massedcompute.com]
- [4. mr-cfd.com](https://mr-cfd.com) [mr-cfd.com]
- [5. NVIDIA H100 80GB PCIe Hands on CFD Simulation](https://servethehome.com) [servethehome.com]
- [6. How does memory bandwidth affect the accuracy of CFD simulations on NVIDIA H100 GPUs compared to other NVIDIA GPUs? - Massed Compute](#) [massedcompute.com]
- [7. trgdatacenters.com](https://trgdatacenters.com) [trgdatacenters.com]
- [8. blog.cadfem.net](https://blog.cadfem.net) [blog.cadfem.net]
- [9. Enabling Simulation-Driven Design with Altair® CFD® and NVIDIA® H100 Tensor Core GPUs](https://altair.com) [altair.com]
- [10. NVIDIA A100 vs H100](https://innovationspace.ansys.com) [innovationspace.ansys.com]
- [11. CPU vs GPU in Ansys Fluent: Exxact, SimuTech, & NVIDIA Collaboration](https://exxactcorp.com) [exxactcorp.com]
- [12. ansys.com](https://ansys.com) [ansys.com]
- [13. How do I optimize OpenFOAM simulations for NVIDIA H100 GPUs, and what are the recommended settings? - Massed Compute](#) [massedcompute.com]
- [14. Fluent GPU Solver Hardware Buying Guide | Ansys Knowledge](https://innovationspace.ansys.com) [innovationspace.ansys.com]
- [15. mdpi.com](https://mdpi.com) [mdpi.com]
- [16. inhalationmag.com](https://inhalationmag.com) [inhalationmag.com]
- [17. hpcwire.com](https://hpcwire.com) [hpcwire.com]

- [18. intuitionlabs.ai \[intuitionlabs.ai\]](https://intuitionlabs.ai)
- To cite this document: BenchChem. [Application Notes and Protocols: Leveraging NVIDIA H100 for Computational Fluid Dynamics Research]. BenchChem, [2026]. [Online PDF]. Available at: [\[https://www.benchchem.com/product/b15585327/docs#application-notes-and-protocols-leveraging-nvidia-h100-for-computational-fluid-dynamics-research\]](https://www.benchchem.com/product/b15585327/docs#application-notes-and-protocols-leveraging-nvidia-h100-for-computational-fluid-dynamics-research)

Disclaimer & Data Validity:

The information provided in this document is for Research Use Only (RUO) and is strictly not intended for diagnostic or therapeutic procedures. While BenchChem strives to provide accurate protocols, we make no warranties, express or implied, regarding the fitness of this product for every specific experimental setup.

Technical Support: The protocols provided are for reference purposes. Unsure if this reagent suits your experiment?

Need Industrial/Bulk Grade? [Request Custom Synthesis Quote](#)

BenchChem

Our mission is to be the trusted global source of essential and advanced chemicals, empowering scientists and researchers to drive progress in science and industry.

Contact

Address: 3281 E Guasti Rd
Ontario, CA 91761, United States
Phone: (601) 213-4426
Email: info@benchchem.com

[Contact our Ph.D. Support Team for a compatibility check](#)