

NVIDIA H100 vs. A100: A Comparative Guide for Scientific Computing

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

Compound of Interest

Compound Name: H100

Cat. No.: B15585327

[Get Quote](#)

For researchers, scientists, and professionals in drug development, the choice of computational hardware is a critical determinant of research velocity and capability. NVIDIA's A100 and **H100** Tensor Core GPUs represent two generations of powerful accelerators that have significantly impacted the landscape of scientific computing. This guide provides an objective comparison of their performance, supported by experimental data, to inform your hardware decisions.

At a Glance: Key Architectural and Performance Differences

The NVIDIA **H100**, based on the Hopper architecture, represents a significant leap forward from the A100, which is built on the Ampere architecture.^[1] This advancement is not merely incremental; the **H100** introduces several new features and substantial performance gains across a range of metrics crucial for scientific workloads.

Table 1: Architectural and Feature Comparison

Feature	NVIDIA A100	NVIDIA H100
GPU Architecture	Ampere[1]	Hopper[1]
CUDA Cores	6,912[2]	14,592[2]
Tensor Cores	3rd Generation[3]	4th Generation[3]
FP64 Performance	9.7 TFLOPS[4]	Up to 60 TFLOPS[5]
FP32 Performance	19.5 TFLOPS[4]	~60 TFLOPS
Memory Type	HBM2e[6]	HBM3[6]
Memory Bandwidth	Up to 2.0 TB/s[7]	Up to 3.35 TB/s[2]
NVLink Interconnect	3rd Generation (600 GB/s)[5]	4th Generation (900 GB/s)[5]
PCIe Support	Gen 4[8]	Gen 5[8]
Multi-Instance GPU (MIG)	1st Generation[3]	2nd Generation[3]

The **H100** boasts more than double the number of CUDA cores, a new generation of Tensor Cores, and significantly higher theoretical peak performance in both double-precision (FP64) and single-precision (FP32) floating-point operations, which are fundamental to scientific simulations.[2][4][5] The upgrade to HBM3 memory and a fourth-generation NVLink interconnect in the **H100** further alleviates potential bottlenecks by providing substantially higher memory bandwidth and faster GPU-to-GPU communication.[5][6][7]

Performance in Scientific Applications

The architectural enhancements of the **H100** translate into tangible performance gains in a variety of scientific computing applications, particularly in fields like molecular dynamics, which are central to drug discovery.

Table 2: Application Performance Benchmarks

Application	Benchmark Dataset	NVIDIA A100 Performance	NVIDIA H100 Performance	Performance Uplift (H100 vs. A100)
GROMACS	STMV (1.06M atoms)	~X ns/day	Up to 2-3x faster	2-3x[5]
NAMD	STMV (1M atoms)	~Y ns/day	Over 2x faster	>2x[9]
LAMMPS	Lennard-Jones (8M atoms)	~Z timesteps/s	~2-3.4x faster	2-3.4x[8]

Note: The performance metrics (ns/day, timesteps/s) are relative and can vary based on the specific system configuration, software versions, and simulation parameters. The performance uplift is a general estimate based on available benchmarks.

For molecular dynamics simulations, the **H100** consistently demonstrates a 2x to over 3x performance improvement compared to the A100.[5][8][9] This acceleration allows researchers to run longer simulations, study larger and more complex biological systems, and obtain results in a fraction of the time.

Experimental Protocols

Reproducible benchmarks are essential for accurate performance assessment. While detailed, end-to-end protocols are not always published in their entirety, the following outlines the general methodologies used in benchmarking these GPUs for scientific applications.



Molecular Dynamics (GROMACS, NAMD, LAMMPS):


- Software Versions: Benchmarks typically utilize recent versions of the simulation packages (e.g., GROMACS 2023, NAMD 3.0).[9][10]
- Input Models: Standard benchmark datasets are often employed to ensure comparability. These include:
 - GROMACS: Satellite Tobacco Mosaic Virus (STMV) with approximately 1 million atoms.

- NAMD: STMV, and Apolipoprotein A1 (APOA1).[11]
- LAMMPS: Lennard-Jones fluid, bead-spring polymer melts, and metallic solids with varying atom counts.[12][13]
- Simulation Parameters:
 - Ensemble: NVT or NPT.
 - Timestep: Typically 2 femtoseconds.
 - Cutoffs: Long-range electrostatics are handled with Particle Mesh Ewald (PME).
 - Constraints: Bonds involving hydrogen atoms are often constrained (e.g., H-bonds in GROMACS).[14]
- Execution: Simulations are run for a set number of steps, and performance is measured in nanoseconds per day (ns/day) or timesteps per second. The initial steps are often excluded from the timing to account for system equilibration.[14]
- Hardware and Software Environment:
 - CPU: High-performance CPUs are used in conjunction with the GPUs.
 - MPI: GPU-aware MPI libraries are used for multi-GPU and multi-node scaling.
 - CUDA Toolkit: The latest versions of the CUDA toolkit are used to leverage new architectural features.[7]

Visualizing the Architectures and Workflows

To better understand the differences and the context of their use, the following diagrams illustrate the architectural advancements and a typical molecular dynamics workflow.





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

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. Exploring the Evolution of GPUs: NVIDIA Hopper vs. Ampere Architectures [blog.paperspace.com]
- 2. vast.ai [vast.ai]
- 3. cudocompute.com [cudocompute.com]
- 4. Nvidia H100 vs A100: A Comparative Analysis [uvation.com]
- 5. What are the key differences between NVIDIA A100 and H100 GPUs for running large-scale molecular dynamics simulations? - Massed Compute [massedcompute.com]
- 6. trgdatacenters.com [trgdatacenters.com]
- 7. H100 vs A100 comparison: Best GPU for LLMs, vision models, and scalable training | Blog — Northflank [northflank.com]
- 8. matsci.org [matsci.org]
- 9. NAMD Performance [ks.uiuc.edu]
- 10. Massively Improved Multi-node NVIDIA GPU Scalability with GROMACS | NVIDIA Technical Blog [developer.nvidia.com]
- 11. Delivering up to 9X the Throughput with NAMD v3 and NVIDIA A100 GPU | NVIDIA Technical Blog [developer.nvidia.com]
- 12. LAMMPS Benchmarks [lammps.org]
- 13. LAMMPS-Bench Molecular Dynamics Benchmark Dataset | Datasets | HyperAI [beta.hyper.ai]
- 14. Creating Faster Molecular Dynamics Simulations with GROMACS 2020 | NVIDIA Technical Blog [developer.nvidia.com]
- To cite this document: BenchChem. [NVIDIA H100 vs. A100: A Comparative Guide for Scientific Computing]. BenchChem, [2026]. [Online PDF]. Available at: [<https://www.benchchem.com/product/b15585327/docs#nvidia-h100-vs-a100-a-comparative-guide-for-scientific-computing>]

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](#)