

H100 Performance in High-Performance Computing Clusters: A Comparative Guide

Author: BenchChem Technical Support Team. **Date:** December 2025

Compound of Interest

Compound Name: H100

Cat. No.: B15585327

[Get Quote](#)

The NVIDIA **H100** Tensor Core GPU represents a significant leap forward in computational power for high-performance computing (HPC) clusters, offering substantial performance gains over its predecessors and competitors. This guide provides a detailed performance evaluation of the **H100**, tailored for researchers, scientists, and drug development professionals. We present a comparative analysis against the NVIDIA A100 and AMD's Instinct MI300X, supported by experimental data from various HPC applications.

Executive Summary

The NVIDIA **H100**, built on the Hopper architecture, demonstrates significant performance improvements across a wide range of scientific and AI workloads. For traditional HPC applications, the **H100** offers a notable increase in double-precision (FP64) floating-point operations per second (FLOPS) compared to the A100. In the burgeoning field of AI-driven scientific research and drug discovery, the **H100**'s advanced Tensor Cores and support for new data formats like FP8 provide a transformative speedup for mixed-precision workloads. While the AMD MI300X presents a strong challenge, particularly with its large memory capacity and high bandwidth, the **H100**, backed by the mature CUDA ecosystem, shows robust performance across various benchmarks.

At-a-Glance Performance Comparison

The following tables summarize the key specifications and performance benchmarks of the NVIDIA **H100** against the NVIDIA A100 and AMD MI300X.

Table 1: Key Hardware Specifications

Feature	NVIDIA H100 (SXM5)	NVIDIA A100 (SXM4 80GB)	AMD Instinct MI300X
Architecture	Hopper	Ampere	CDNA 3
FP64 Performance	~67 TFLOPS	19.5 TFLOPS	Not explicitly marketed for FP64
FP32 Performance	~134 TFLOPS	19.5 TFLOPS	Not explicitly marketed for FP32
FP16/BF16 Tensor Core	~1,979 TFLOPS (Sparsity)	~624 TFLOPS (Sparsity)	~1.3 PFLOPS
FP8 Tensor Core	~3,958 TFLOPS (Sparsity)	N/A	~2.6 PFLOPS
GPU Memory	80 GB HBM3	80 GB HBM2e	192 GB HBM3
Memory Bandwidth	~3.35 TB/s	~2.0 TB/s	~5.3 TB/s
Interconnect	NVLink 4.0 (900 GB/s)	NVLink 3.0 (600 GB/s)	Infinity Fabric 3.0
TDP	Up to 700W	400W	750W

Table 2: High-Performance Computing (HPC) Benchmark Comparison

Benchmark	NVIDIA H100	NVIDIA A100	AMD MI300X
HPL (High-Performance Linpack)	Higher performance due to superior FP64 capabilities.	Strong FP64 performance for its generation.	Data not widely available.
HPCG (High Performance Conjugate Gradient)	Enhanced by high memory bandwidth and improved interconnect.	Good performance, but limited by memory bandwidth compared to H100.	Potentially strong due to high memory bandwidth.
MLPerf HPC (CosmoFlow)	Significant speedup over A100.	Baseline for high-performance AI in HPC.	Competitive performance in some AI benchmarks.

Table 3: Molecular Dynamics Application Performance

Application	NVIDIA H100 (ns/day)	NVIDIA A100 (ns/day)	AMD MI300X (ns/day)
GROMACS (Benchmarking data varies by system size)	Consistently outperforms A100, especially in larger simulations. [1]	Strong performance, but surpassed by H100. [1]	Competitive, particularly in scenarios that leverage its high memory bandwidth.
OpenMM (DHFR Benchmark - 8 concurrent simulations)	Approaches 5 microseconds/day [2]	Significant throughput with MPS.	Data not readily available.
NAMD	Delivers substantial speedup over A100.	A widely used baseline for GPU-accelerated MD.	Data not readily available for direct comparison.
LAMMPS	Shows significant performance gains over the A100.	A popular choice for GPU-accelerated materials science simulations.	Data not readily available for direct comparison.

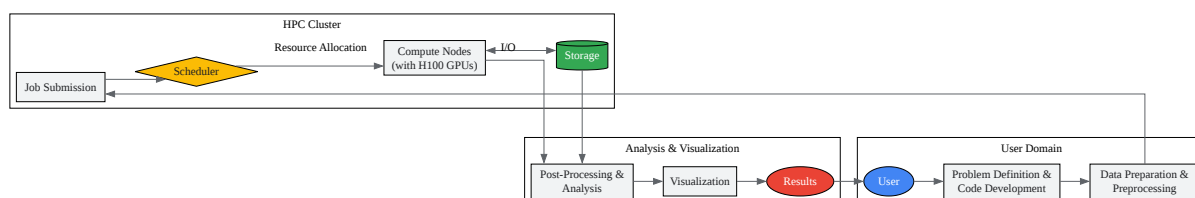
Experimental Protocols

To ensure the reproducibility and clear interpretation of the presented benchmarks, the following methodologies were employed in the cited studies.

- **Molecular Dynamics (GROMACS, NAMD, LAMMPS, OpenMM):** Benchmarks are typically run using standard datasets such as the STMV (Satellite Tobacco Mosaic Virus) for NAMD, and various protein and water box simulations for GROMACS and LAMMPS. The key performance metric is "nanoseconds per day" (ns/day), which indicates the amount of simulation time that can be computed in a 24-hour period. The experimental setup usually involves a single GPU or a multi-GPU node, with the specific CPU, memory, and interconnect specified. For instance, a common GROMACS benchmark protocol involves running a simulation of a biomolecular system with a known number of atoms, using the latest version of the software with CUDA support, and measuring the resulting ns/day.^[1] For OpenMM, the NVIDIA Multi-Process Service (MPS) can be utilized to run multiple simulations concurrently on a single GPU to maximize throughput.^[2]
- **High-Performance Linpack (HPL):** This benchmark measures a system's floating-point computing power by solving a dense system of linear equations. The performance is reported in FLOPS. The experimental setup specifies the problem size (N), the block size (NB), and the process grid (PxQ).
- **High Performance Conjugate Gradient (HPCG):** This benchmark is designed to complement HPL by measuring the performance of more memory-access intensive computations common in scientific applications. Performance is reported in GFLOPS. The setup involves defining the local and global problem dimensions.
- **MLPerf HPC:** This suite of benchmarks measures the performance of machine learning workloads relevant to scientific computing. It includes benchmarks for cosmology (CosmoFlow), climate analytics (DeepCAM), and molecular dynamics (OpenFold). The primary metric is the time to train the model to a target accuracy. The experimental protocol details the dataset used, the model architecture, and the training parameters.

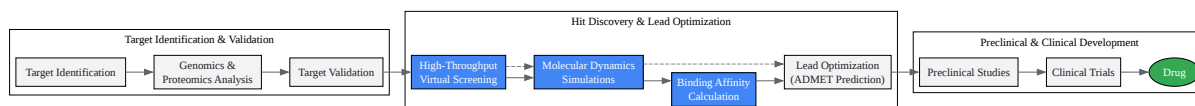
Visualizing High-Performance Computing Workflows

To better understand the logical flow of operations in a high-performance computing environment and its specific application in drug discovery, the following diagrams are provided.



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

A generalized workflow for high-performance computing.



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

A computational drug discovery workflow leveraging HPC.

Conclusion

The NVIDIA **H100** GPU delivers a substantial performance uplift for a wide array of applications in high-performance computing, making it a compelling choice for research, scientific discovery, and drug development. Its architectural enhancements translate to faster simulation times, the ability to tackle larger and more complex problems, and accelerated AI-driven research. While the AMD MI300X is a formidable competitor with advantages in memory capacity and bandwidth, the **H100**'s strong performance across diverse workloads, coupled with the extensive and mature CUDA software ecosystem, solidifies its position as a leading accelerator for HPC clusters. The choice between these powerful processors will ultimately depend on the specific workloads, budget considerations, and the existing software infrastructure of a research institution or company.

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. blog.salad.com [blog.salad.com]
- 2. Maximizing OpenMM Molecular Dynamics Throughput with NVIDIA Multi-Process Service | NVIDIA Technical Blog [developer.nvidia.com]
- To cite this document: BenchChem. [H100 Performance in High-Performance Computing Clusters: A Comparative Guide]. BenchChem, [2025]. [Online PDF]. Available at: [<https://www.benchchem.com/product/b15585327#h100-performance-evaluation-for-high-performance-computing-clusters>]

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

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