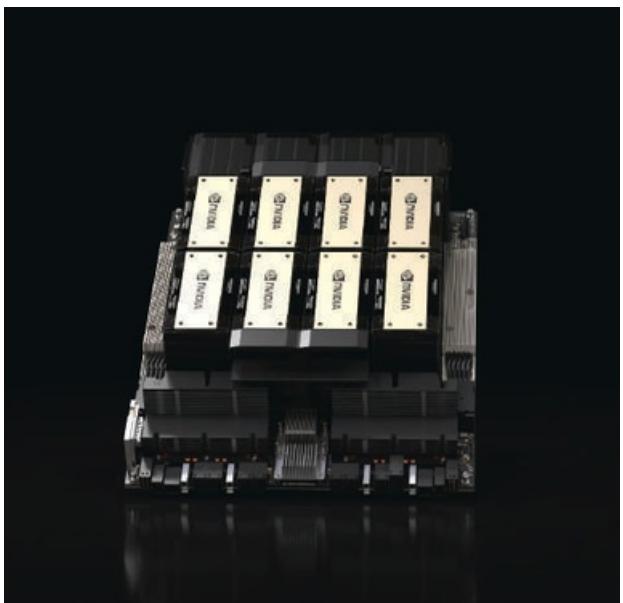




# NVIDIA H200 Tensor Core GPU

Supercharging AI and HPC workloads.



## Higher Performance With Larger, Faster Memory

The NVIDIA H200 Tensor Core GPU supercharges generative AI and high-performance computing (HPC) workloads with game-changing performance and memory capabilities.

Based on the **NVIDIA Hopper™ architecture**, the NVIDIA H200 is the first GPU to offer 141 gigabytes (GB) of HBM3e memory at 4.8 terabytes per second (TB/s)—that's nearly double the capacity of the **NVIDIA H100 Tensor Core GPU** with 1.4X more memory bandwidth. The H200's larger and faster memory accelerates generative AI and large language models, while advancing scientific computing for HPC workloads with better energy efficiency and lower total cost of ownership.

## Key Features

- 141GB of HBM3e GPU memory
- 4.8TB/s of memory bandwidth
- 4 petaFLOPS of FP8 performance
- 2X LLM inference performance
- 110X HPC performance

## Unlock Insights With High-Performance LLM Inference

In the ever-evolving landscape of AI, businesses rely on large language models to address a diverse range of inference needs. An **AI inference** accelerator must deliver the highest throughput at the lowest TCO when deployed at scale for a massive user base.

The H200 doubles inference performance compared to H100 GPUs when handling large language models such as Llama2 70B.

### Up to 2X the LLM Inference Performance

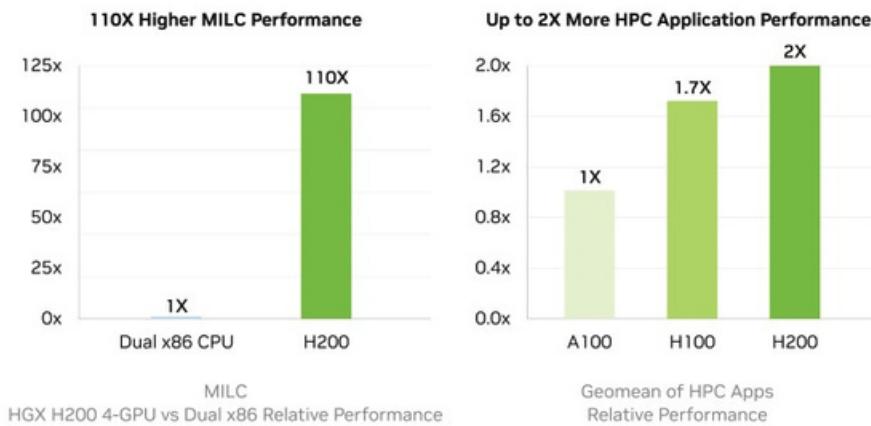


Preliminary specifications. May be subject to change.

Llama2 13B: ISL 128, OSL 2K | Throughput | H100 SXM 1x GPU BS 64 | H200 SXM 1x GPU BS 128  
 GPT-3 175B: ISL 80, OSL 200 | x8 H100 SXM GPUs BS 64 | x8 H200 SXM GPUs BS 128  
 Llama2 70B: ISL 2K, OSL 128 | Throughput | H100 SXM 1x GPU BS 8 | H200 SXM 1x GPU BS 32.

## Supercharge High-Performance Computing

Memory bandwidth is crucial for HPC applications, as it enables faster data transfer and reduces complex processing bottlenecks. For memory-intensive HPC applications like simulations, scientific research, and artificial intelligence, the H200's higher memory bandwidth ensures that data can be accessed and manipulated efficiently, leading to 110X faster time to results.



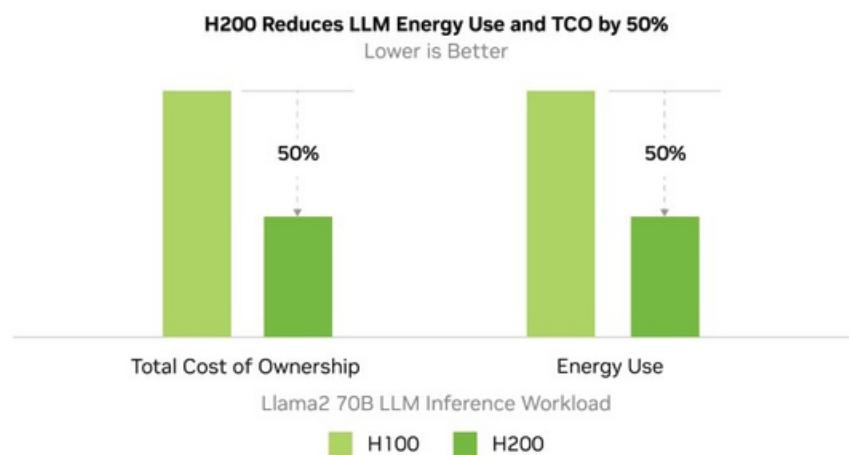
Preliminary specifications. May be subject to change.

HPC MILC: dataset NERSC Apex Medium | HGX H200 4-GPU | dual Sapphire Rapids 8480

HPC Apps: CP2K: dataset H2O-32-RI-dRPA-96points | GROMACS: dataset STMV | ICON: dataset r2b5 | MILC: dataset NERSC Apex Medium | Chroma: dataset HMC Medium | Quantum Espresso: dataset AUSURF112 | 1x H100 SXM | 1x H200 SXM.

## Reduce Energy and TCO

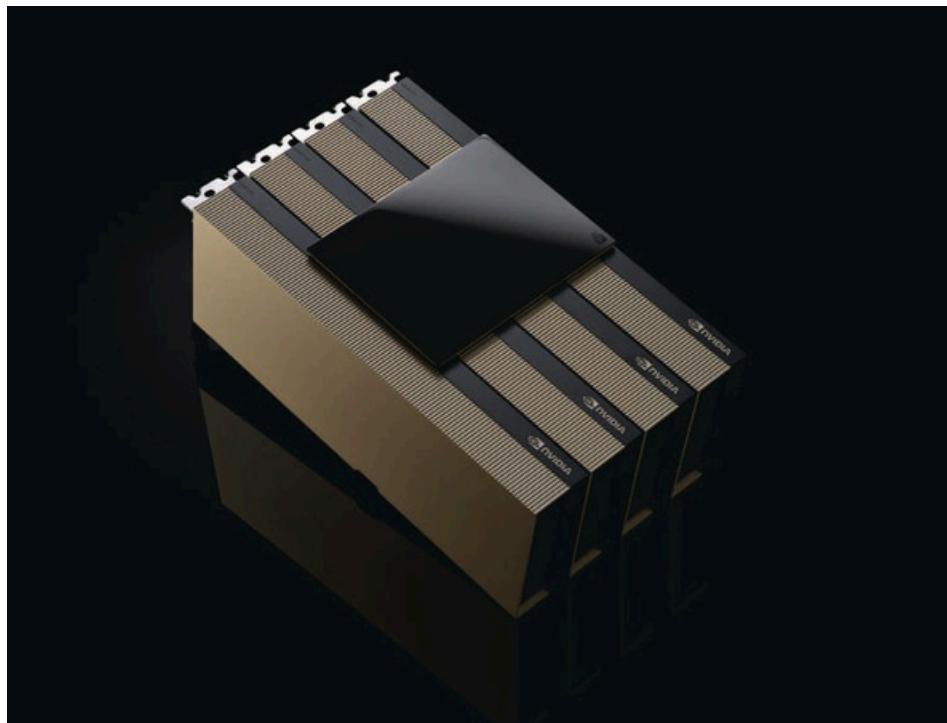
With the introduction of H200, energy efficiency and TCO reach new levels. This cutting-edge technology offers unparalleled performance, all within the same power profile as the **H100 Tensor Core GPU**. AI factories and supercomputing systems that are not only faster but also more eco-friendly deliver an economic edge that propels the AI and scientific communities forward.



Preliminary specifications. May be subject to change.

LLama2 70B: ISL 2K, OSL 128 | Throughput | H100 SXM 1x GPU BS 8 | H200 SXM 1x GPU BS 32

## AI Acceleration for Mainstream Enterprise Servers With H200 NVL



NVIDIA H200 NVL is ideal for lower-power, air-cooled enterprise rack designs that require flexible configurations, delivering acceleration for every AI and HPC workload regardless of size. With up to four GPUs connected by **NVIDIA NVLink™** and a 1.5X memory increase, large language model (LLM) inference can be accelerated up to 1.7X and HPC applications achieve up to 1.3X more performance over the H100 NVL.

### Enterprise-Ready: AI Software Streamlines Development and Deployment

NVIDIA H200 NVL comes with a five-year **NVIDIA AI Enterprise** subscription and simplifies the way you build an enterprise AI-ready platform. H200 accelerates AI development and deployment for production-ready generative AI solutions, including computer vision, speech AI, retrieval augmented generation (RAG), and more. NVIDIA AI Enterprise includes **NVIDIA NIM™**, a set of easy-to-use microservices designed to speed up enterprise generative AI deployment. Together, deployments have enterprise-grade security, manageability, stability, and support. This results in performance-optimized AI solutions that deliver faster business value and actionable insights.

## Technical Specifications

	<b>H200 SXM1</b> 34 TFLOPS 67	<b>H200 NVL1</b> 30 TFLOPS 60 TFLOPS 60
<b>FP64</b>	TFLOPS 67 TFLOPS 989	TFLOPS 835 TFLOPS 1,671 TFLOPS 1,671
<b>FP64 Tensor Core</b>	TFLOPS 1,979 TFLOPS	TFLOPS 3,341 TFLOPS 3,341 TFLOPS 141GB
<b>FP32</b>	1,979 TFLOPS 3,958	4.8TB/s 7 NVDEC
<b>TF32 Tensor Core2</b>	TFLOPS 3,958 TFLOPS	7 JPEG
<b>BFLOAT16 Tensor Core2</b>	141GB 4.8TB/s 7 NVDEC	Supported
<b>FP16 Tensor Core2</b>	7 JPEG	Up to 600W (configurable)
<b>FP8 Tensor Core2</b>	Supported	Up to 7 MIGs @16.5GB each
<b>INT8 Tensor Core2</b>	Up to 700W (configurable)	PCIe
<b>GPU Memory</b>	Up to 7 MIGs @18GB each	Dual-slot air-cooled 2- or 4-way NVIDIA
<b>GPU Memory Bandwidth</b>	SXM	NVLink bridge: 900GB/s per GPU
<b>Decoders</b>		PCIe Gen5: 128GB/s NVIDIA MGX™ H200 NVL partner and
<b>Confidential Computing</b>		NVIDIA-Certified Systems with up to 8 GPUs
<b>Max Thermal Design Power (TDP)</b>		Included
<b>Multi-Instance GPUs</b>		
<b>Form Factor</b>		
<b>Interconnect</b>	NVIDIA NVLink: 900GB/s PCIe Gen5: 128GB/s	
<b>Server Options</b>	NVIDIA HGX™ H200 partner and NVIDIA-Certified Systems™ with 4 or 8 GPUs	
<b>NVIDIA AI Enterprise</b>	Add-on	

1. Preliminary specifications. May be subject to change.

2. With sparsity.

## Ready to Get Started?

To learn more about the NVIDIA H200 Tensor Core GPU,

visit [nvidia.com/h200](https://nvidia.com/h200)

© 2024 NVIDIA Corporation and affiliates. All rights reserved. NVIDIA, the NVIDIA logo, HGX, Hopper, MGX, NIM, NVIDIA-Certified Systems, and NVLink are trademarks and/or registered trademarks of NVIDIA Corporation and affiliates in the U.S. and other countries. Other company and product names may be trademarks of the respective owners with which they are associated. 3512650. NOV24

