




NVIDIA H100 TENSOR CORE GPU

Unprecedented performance, scalability, and security for every data center.



Take an order-of-magnitude leap in accelerated computing.

The NVIDIA H100 Tensor Core GPU delivers unprecedented performance, scalability, and security for every workload. With NVIDIA® NVLink®, two H100 PCIe GPUs can be connected to accelerate demanding compute workloads, while the dedicated Transformer Engine supports large parameter language models. H100 uses breakthrough innovations in the NVIDIA Hopper™ architecture to deliver industry-leading conversational AI, speeding up large language models by 30x over the previous generation.

Ready for enterprise AI?

NVIDIA H100 Tensor Core GPUs for mainstream servers come with a five-year software subscription, including enterprise support, to the NVIDIA AI Enterprise software suite, simplifying AI adoption with the highest performance. This ensures organizations have access to the AI frameworks and tools they need to build H100-accelerated AI workflows such as AI chatbots, recommendation engines, vision AI, and more. Access the NVIDIA AI Enterprise software subscription and related support benefits for the NVIDIA H100 [here](#).

Securely accelerate workloads from enterprise to exascale.

NVIDIA H100 GPUs feature fourth-generation Tensor Cores and the Transformer Engine with FP8 precision, further extending NVIDIA's market-leading AI leadership with up to 9x faster training and an incredible 30x inference speedup on large language models. For high-performance computing (HPC) applications, H100 triples the floating-point operations per second (FLOPS) of FP64 and adds dynamic programming (DPX) instructions to deliver up to 7x higher performance. With second-generation Multi-Instance GPU (MIG), built-in NVIDIA confidential computing, and NVIDIA NVLink, H100 securely accelerates all workloads for every data center from enterprise to exascale.


SPECIFICATIONS

PNY Part Number	NVH100TCGPU-KIT
FP64	26 TFLOPS
FP64 Tensor Core	51 TFLOPS
FP32	51 TFLOPS
TF32 Tensor Core	51 TFLOPS*
BFLOAT16 Tensor Core	1,513 TFLOPS*
FP16 Tensor Core	1,513 TFLOPS*
FP8 Tensor Core	3,026 TFLOPS*
INT8 Tensor Core	3,026 TOPS*
GPU memory	80GB HBM2e
GPU memory bandwidth	2TB/s
Decoders	7 NVDEC 7 JPEG
Max thermal design power (TDP)	300-350W (configurable)
Multi-Instance GPUs	Up to 7 MIGS at 10GB each
Form factor	PCIe Dual-Slot Passive Fansink
Interconnect	NVLink: 600GB/s, PCIe Gen5: 128GB/s
Server options	NVIDIA and Partner-Certified Systems with 1-8 GPUs
NVIDIA AI Enterprise	Included

* Shown with sparsity. Specifications 1/2 lower without sparsity.

Up to 9x Higher AI Training on Largest Models


Mixture of Experts (395 Billion Parameters)



Projected performance subject to change. Training Mixture of Experts (MoE) Transformer Switch-XXL variant with 395B parameters on 1T token dataset | A100 cluster: HDR IB network | H100 cluster: NVLink, NDR IB


Up to 30x Higher AI Inference Performance on Largest Models

Megatron Chatbot Inference (530 Billion Parameters)



Projected performance subject to change. Inference on Megatron 530B parameter model chatbot for input sequence length=128, output sequence length=20 | A100 cluster: HDR IB network | H100 cluster: NDR IB network for 16 H100 configurations | 32 A100 vs 16 H100 for 1 and 1.5 sec | 16 A100 vs 8 H100 for 2 sec

Up to 7x Higher Performance for HPC Applications



Projected performance subject to change. 3D FFT (4K^3) throughput | A100 cluster: HDR IB network | H100 cluster: NVLink, NDR IB | Genome Sequencing (Smith-Waterman) | 1 A100 | 1 H100


Explore the technology breakthroughs of NVIDIA Hopper.



NVIDIA H100 Tensor Core GPU

Built with 80 billion transistors using a cutting-edge TSMC 4N process custom tailored for


NVIDIA's accelerated compute needs, H100 is the world's most advanced chip ever built. It features major advances to accelerate AI, HPC, memory bandwidth, interconnect, and communication at data center scale.



Transformer Engine

The Transformer Engine uses software and Hopper Tensor Core technology designed to accelerate training for models

built from the world's most important AI model building block, the transformer. Hopper Tensor Cores can apply mixed FP8 and FP16 precisions to dramatically accelerate AI calculations for transformers.



NVLink Switch System

NVLink enables the scaling of dual H100 input/output (IO) at 600 gigabytes per second (GB/s) bidirectional per GPU, over 7x

the bandwidth of PCIe Gen5, and also delivers 9x higher bandwidth than InfiniBand HDR on the NVIDIA Ampere architecture.



NVIDIA Confidential Computing

NVIDIA Confidential Computing is a built-in security feature of Hopper that makes NVIDIA H100


the world's first accelerator with confidential computing capabilities. Users can protect the confidentiality and integrity of their data and applications in use while accessing the unsurpassed acceleration of H100 GPUs.



Second-Generation Multi-Instance GPU (MIG)

The Hopper architecture's second-generation MIG supports multi-tenant,

multi-user configurations in virtualized environments, securely partitioning the GPU into isolated, right-size instances to maximize quality of service (QoS) for 7x more secured tenants.



DPX Instructions

Hopper's DPX instructions accelerate dynamic programming algorithms by 40x compared to CPUs and


7x compared to NVIDIA Ampere architecture GPUs. This leads to dramatically faster times in disease diagnosis, real-time routing optimizations, and graph analytics.

Accelerate every workload, everywhere.

The NVIDIA H100 is an integral part of the NVIDIA data center platform. Built for AI, HPC, and data analytics, the platform accelerates over 3,000 applications, and is available everywhere from data center to edge, delivering both dramatic performance gains and cost-saving opportunities.

Deploy H100 with the NVIDIA AI platform.

NVIDIA AI is the end-to-end open platform for production AI built on NVIDIA H100 GPUs. It includes NVIDIA accelerated computing infrastructure, a software stack for infrastructure optimization and AI development and deployment, and application workflows to speed time to market. Experience NVIDIA AI and **NVIDIA H100 on NVIDIA LaunchPad** through free hands-on labs.



Ready to Get Started?

To learn more about the NVIDIA H100 Tensor Core GPU, visit: www.pny.com/h100