Hewlett Packard Enterprise

REDUCE COMPLEXITY WITH HPE APOLLO SYSTEMS AND NVIDIA GPUS FOR AI WORKLOADS

Accelerate AI solutions for training and inference

Technology is transforming nearly every human and business process, from driving business growth, to translating documents in real time, to enhancing decision-making in areas such as financial services and scientific research. Spearheading these technological advancements, artificial intelligence (AI) is demanding a new breed of performance-accelerated machines that can solve highly complex problems quickly, while simplifying IT management and reducing time to insight.

HPE helps you accelerate value from your data, using AI to create new experiences, drive smarter operations, and deliver breakthrough innovation.

HPE APOLLO AI STARTER KIT WITH NVIDIA® T4 GPUS

The HPE Apollo 2000 Gen10 System offers density-optimized, scale-out computing for enterprises looking to get started with machine learning. Powered by NVIDIA T4 enterprise GPUs and certified to run NVIDIA GPU Cloud (NGC) pre-build containers for high-performance computing (HPC) and machine learning, HPE Apollo 2000 Gen10 System takes training and inference workloads to new heights of performance, efficiency, and security. The HPE Apollo AI Starter Kit is designed to provide everything you need to get started with serious AI development and deployment.

HPE APOLLO AI TRAINING KIT WITH NVIDIA V100 TENSOR CORE GPUS

HPE Apollo 6500 Gen10 System is a platform specially engineered for machine learning and provides unprecedented performance with up to

eight NVIDIA V100 Tensor Core GPUs utilizing the NVIDIA NVLink interconnect. As an NGC-Ready validated server, you can utilize the NGC hub to immediately download GPU-optimized software for deep learning, machine learning, and HPC that handles the infrastructure. This helps data scientists, developers, and researchers to focus on building solutions, gathering insights, and delivering business value.[1]

HPE APOLLO GEN10 SYSTEMS: BUILT FOR HPC, ENHANCED FOR AI

HPE takes training and inference workloads to new heights of performance, efficiency, and security by offering bundles that include essential building blocks that can scale. By applying these unique configurations for training and inference, you benefit from:

- HPE's industry-leading server technology that delivers unprecedented performance with reliability, availability, and services (RAS) features

- A complete solution, including the right mix of NVIDIA GPUs and Intel® Xeon® CPUs at a discounted price

- NVIDIA Deep Learning Institute (DLI) training credits for insight into deep learning techniques and best practices

- HPE Pointnext Services and support

- 3-year parts, labor, and on-site support warranty

- Simple system management

- Essential firmware anchored by the HPE iLO 5 chip by Silicon Root of Trust to create an immutable fingerprint that verifies the firmware code is valid, so the server won't boot with compromised firmware

[1] nvidia.com/en-us/gpu-cloud/

- Superior performance per dollar, ease of serviceability, and zoned cooling with GPUs

- Broad choice of tier one OS supported

- Added security and smart remote functionality with the HPE iLO 5 Advanced License

VALIDATED SOLUTION ELEMENTS FOR THE AI STARTER KIT WITH HPE APOLLO 2000

HPE Apollo 2000 Gen10 System can be ordered with either four or eight NVIDIA T4 16 GB GPUs. The bundle consists of:

- HPE Apollo 2000 Gen10 Server (HPE ProLiant XL190r)

- Intel® Xeon® Gold 6230 (2.1 GHz/20-core/125W) processor

- 192 Gb RAM/384 Gb RAM

- 4.8 TB storage/9 or 6 TB storage

- 10/25GbE ports

- HPE iLO Advanced 1-server License with 3-year support on iLO licensed features

- Six NVIDIA DLI online course licenses

- HPE Foundation Care services

The chassis with a single node gives room to add a second node and scales up to eight GPUs per chassis. This helps you meet the demanding data read/write requirements on the storage and data management components of HPC analytics and AI environments.

VALIDATED SOLUTION ELEMENTS FOR THE AI TRAINING KIT WITH HPE APOLLO 6500 GEN10 SYSTEM

Built for machine and deep learning workloads, HPE Apollo 6500 Gen10 System comes with either four or eight NVIDIA V100 Tensor Core Accelerators. The bundle consists of:

- HPE Apollo 6500 Gen10 Server

- Intel Xeon Gold 6230 (2.1 GHz/20-core/125W)

- NVIDIA V100 32 GB SXM2 GPUs

- NVIDIA NVLink 2.0 GPU Cross-connect

- 768 Gb RAM

- 8.64 TB Storage

- 10/25GbE ports

- Two InfiniBand EDR/Ethernet 100 Gb ports

- HPE iLO Advanced 1-server license with 3-year support on iLO licensed features

- Six NVIDIA DLI online course licenses

- HPE Foundation Care services

A POWERFUL SOLUTION SUPPORTED BY GLOBAL SERVICES

HPE Proactive Care gives customers an enhanced call experience. When your products are connected to Hewlett Packard Enterprise, HPE Proactive Care helps prevent problems and maintains IT stability by utilizing personalized proactive reports with recommendations and advice. This service combines both reactive support when there is a problem with an enhanced call experience and start to finish case management with proactive reporting and advice. Additionally, HPE Pointnext Services provides a comprehensive portfolio including Advisory and Transformational, Professional, and Operational Services to help accelerate your digital transformation.

LEARN MORE AT
hpe.com/us/en/solutions/artificial-intelligence.html

Make the right purchase decision. Contact our presales specialists.

Chat **Email** **Call**

Get updates

Hewlett Packard Enterprise