Hewlett Packard Enterprise

Solution overview

HPE Apollo 6500 Gen10 System

Enterprise platform for accelerated computing

Reliable enterprise deep learning and **high-performance computing (HPC) platform** with industry-leading accelerators delivers exceptional performance for faster intelligence.

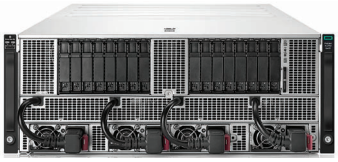

Demystify deep learning for faster intelligence and achieve greater researcher productivity with the HPE Apollo 6500 Gen10 System.

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

Today's organizations are leveraging **deep learning**—a powerful component of AI—to analyze data and derive actionable intelligence at lightning speeds. Backed by HPC operations, deep learning applications utilize predictive analytics to forecast future activity, behaviors, and trends, as well as use Big Data analytics to disclose and resolve hidden problems in data.

Capitalizing on the full range of AI and deep learning capabilities requires purpose-built computers capable of learning freely, reasoning, and determining the most appropriate course of action in real time. The HPE Apollo 6500 Gen10 System best addresses the most important step of training the deep learning model.

Enterprise customers with modernized applications are looking to reap the benefits of accelerated computing to expand their use of HPC power to speed research and bring new products and services to market faster. With the HPE Apollo 6500 Gen10 System, enterprise customers receive:

- Unprecedented performance
- Economical AI and deep learning
- Rock-solid, enterprise-level reliability, availability, serviceability (RAS) features
- Easier system management
- Flexibility to support a wide range of workloads, including deep learning and typical HPC workloads of complex simulation and modeling

Greater performance, flexibility, resiliency, and more

With eight high-performance graphical processing units (GPUs) per server, the HPE Apollo 6500 Gen10 System provides superior performance per dollar for GPU-intensive workloads—delivering up to 125 TFLOPs[2] single-precision compute. Purpose-built for accelerated computing, this platform features both PCIe and NVLink GPU interconnects, providing the flexibility to suit a wide variety of requirements.

The NVLink 2.0 GPU interconnect is particularly useful for deep learning workloads, characterized by heavy GPU-to-GPU communications. High-bandwidth, low-latency networking adapters (up to four high-speed Ethernet, Intel® Omni-Path Architecture, InfiniBand Enhanced Data Rate [EDR], and future InfiniBand HDR per server) are tightly coupled with the GPU accelerators, which allows the system to take full advantage of the network bandwidth.

Leveraging the groundbreaking advancements of **HPE ProLiant** Gen10 servers, the HPE Apollo 6500 Gen10 System offers exceptional:

- **Agility**—To accelerate your business with HPE's highest number of GPUs per server
- **Security**—To provide the ultimate in IT protection by running the world's most secure industry-standard servers[1]
- **Economic control**—To reduce operational costs by leveraging a pay-as-you-go consumption-based model

[1] Based on external firm conducting cybersecurity penetration testing of a range of server products from a range of manufacturers, May 2017.

[2] Theoretical peak performance.

The HPE Apollo 6500 Gen10 System improves HPC and deep learning researcher productivity with:

- **Accelerated performance for GPU-intensive workloads**
  - Leading accelerator technology, using NVLink 2.0 for dedicated GPU-to-GPU communication
  - HPE's highest number of GPUs per server
  - Powerful host, with high-speed/low-latency network fabric adapters, NVMe drives, and high-speed DDR4 SmartMemory
  - Dependable performance, with power and cooling designed around 350 watt accelerators and consistent signal integrity for reliable operations
- **Flexibility for both HPC and deep learning environments**
  - Accelerator choice, either NVLink 2.0 for maximum bandwidth or PCIe for traditional GPU support
  - Flexible support, enterprise options, and choice of Ubuntu® or enterprise-grade Linux® distribution (Red Hat®, SUSE®, CentOS) to match your business strategy
  - Accelerator topologies to suit varying workloads, such as an efficient hybrid cube mesh for NVLink and 4:1 or 8:1 GPU:CPU flexibility in PCIe
  - Broad **storage options**, with up to 16 front-accessible storage devices—SAS/SATA solid-state drives (SSDs) with up to 4 NVMe drives
- **Resiliency, security, and simplicity for lower TCO**
  - Easier serviceability and upgrades, with easy-access modular design and rear-cabled fabrics
  - Optimized power and cooling, with 2 + 2 power for resiliency
  - Efficient system management with **HPE iLO 5** (Integrated Lights-Out) for lower TCO and firmware-level rock-solid security
  - All-in-one design with integrated power supplies simplifies deployment in a standard 1 m deep rack

Deep learning and HPE Apollo 6500—the perfect match

Making AI real for wider application areas, deep learning uses high-performance computers to identify patterns and relationships within massive amounts of data. To date, traditional HPC systems were unable to keep pace with the unprecedented demands required by deep learning. As such, the benefits of deep learning remained virtually untapped.

Today, however, the HPE Apollo 6500 Gen10 System delivers the necessary compute capacity to enable organizations of all sizes to realize the true benefits of deep learning. With the HPE Apollo 6500 Gen10 System, organizations can put deep learning into action by:

- Identifying vehicles, pedestrians, and landmarks; for example, autonomous vehicles
- Monitoring oil field drilling rigs to prevent disasters
- Recognizing images
- Recognizing speech and translating
- Processing natural languages
- Designing drugs
- Completing bioinformatics

Technical specifications

HPE ProLiant XL270d Gen10 Server

Graphical processing units	Up to eight GPUs per server, supporting up to 350 watt GPUs, having fast dedicated GPU-to-GPU communication option with NVLink 2.0 or standard PCIe
Supported accelerators	NVIDIA® Tesla V100, P100, P40 in both NVLink and PCIe configurations
Adapters	Four high-speed fabric adapters (Ethernet, Intel Omni-Path Architecture, InfiniBand EDR, and future InfiniBand HDR)
Processor	Intel® Xeon® Scalable processors (8100 and 6100); up to 165 watt, having up to 28 cores or up to 3 GHz frequency
Memory	24 of 2666 MHz DDR4 SmartMemory DIMMs, providing 3 TB of memory (24 x 128 GB)
Storage	Up to 16 SFF devices: hot-plug SAS/SATA SSD with up to 4 NVMe SSD front-accessible storage
Network options	• Four 1GbE single-port RJ-45 • Four x 16 PCI Express slots supporting InfiniBand EDR or Intel OPA or high-speed Ethernet, up to 100 Gbps
I/O slots	• 1 x 16 FHHL PCIe Gen3 System Board module • 4 x 16 PCIe low-profile Gen3 from GPU module
Power supplies	Up to four 2200W 80 Plus Platinum (2 + 2)
Rack	Designed for standard 1075 mm deep racks

HPE demystifies deep learning for faster intelligence

To support enterprise customers at various stages along the AI/deep learning adoption continuum, HPE created the Deep Learning Cookbook for Starters, complete with:

- Recipes for deep learning workloads
- Proven blueprints and reference configurations for customers integrating deep learning workloads in their organizations

HPE financing for HPE Apollo 6500 Gen10 System

A critical component of success is having access to technology on terms that align to your business requirements and budget. **HPE Financial Services** is uniquely positioned to deliver a broad portfolio of flexible investment and transition solutions designed to accelerate your move to the data center of the future.

Get started today

Leveraging the power of AI and deep learning to solve complex social, scientific, and engineering problems starts by partnering with Hewlett Packard Enterprise. Contact your HPE representative today. Find out how, together we can put the HPE Apollo 6500 Gen10 System to work for your organization.

Learn more at

hpe.com/us/en/solutions/hpc-high-performance-computing/deep-learning.html

hpe.com/us/en/solutions/hpc-high-performance-computing.html

HPE Pointnext

- Providing a comprehensive portfolio of services to help accelerate digital transformation. Focusing on creative configurations with flawless on-time implementation and on-budget execution. Following innovative approaches—**HPE GreenLake** and **HPE Datacenter Care**—to keep businesses at peak performance.

HPE Artificial Intelligence Transformation Workshop

- A one-day workshop that helps you clarify AI concepts, identify priority use cases for your business, and create a high-level plan that defines your steps forward. Senior **HPE Pointnext** experts with AI, data, and analytics expertise use informative visual displays in an interactive session to share their knowledge with you.
- Additional offerings include proof-of-value services based on industry/machine learning use cases and full implementation.

Sign up for updates

a00041561ENW, March 2018