

# Reduce complexity with HPE Apollo and NVIDIA GPUs for AI workloads

## An ideal training solution for machine learning

The Artificial Intelligence (AI) landscape is complicated and businesses are realizing that massive influx of data is only useful if they can glean insights from it. HPE will help you accelerate value from all of your data, using AI to create new experiences, drive smarter operations, and breakthrough innovation.

and researchers can focus on building solutions, gathering insights, and delivering business value.<sup>1</sup>






### HPE Apollo AI training bundle with NVIDIA® GPUs

The HPE Apollo 6500 Gen10 System is an ideal machine learning platform providing unprecedented performance with up to eight NVIDIA Tesla® V100 GPUs utilizing NVLink interconnect. As an NVIDIA GPU Cloud (NGC) ready server supporting NVIDIA GPUs, you can utilize the NGC hub to immediately download GPU-optimized software for deep learning, machine learning, and HPC that takes care of all the plumbing so data scientists, developers,

### A better bundle built for training models

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 like financial services and scientific research. Spearheading these technological advancements, AI is introducing a new breed of intelligent machines that can solve highly complex problems quickly, while simplifying IT management and reducing time to insight. Start accelerating your AI workloads today with the HPE Apollo 6500 Gen10 system. The following are some key use cases where the HPE Apollo 6500 system excels.

**Table 1.** Key uses cases for the HPE Apollo 6500 Gen10 bundle

 Financial services	 Healthcare	 Telecom	 Marketing and retail	 Manufacturing
<ul style="list-style-type: none"> <li>Know your customers (KYC)</li> <li>Anti-money laundering (AML)</li> <li>Transaction fraud</li> <li>Collusion fraud</li> <li>Real-time targeting</li> <li>Credit risk scoring</li> <li>In-context promotion</li> <li>Deposit fraud</li> <li>Customer churn prediction</li> <li>Auto loan</li> </ul>	<ul style="list-style-type: none"> <li>Early cancer detection</li> <li>Product recommendations</li> <li>Personalized prescription matching</li> <li>Medical claim fraud detection</li> <li>Flu season prediction</li> <li>Drug discovery</li> <li>ER and hospital management</li> <li>Remote patient monitoring</li> <li>Medical test predictions</li> </ul>	<ul style="list-style-type: none"> <li>Predictive maintenance</li> <li>Avoidable truck rolls</li> <li>Customer churn prediction</li> <li>Improved customer viewing experience</li> <li>Master data management</li> <li>In-context promotions</li> <li>Intelligent and placements</li> <li>Personalized program recommendations</li> </ul>	<ul style="list-style-type: none"> <li>Funnel predictions</li> <li>Personalized ads</li> <li>Credit scoring</li> <li>Fraud detection</li> <li>Next best offer</li> <li>Next best customer</li> <li>Smart profiling</li> <li>Prediction</li> <li>Customer recommendations</li> </ul>	<ul style="list-style-type: none"> <li>Design exploration</li> <li>Simulation/physical verification testing</li> <li>Hybrid twin</li> <li>Video/predictive analytics</li> <li>Logistics optimization</li> <li>Visual inspection</li> <li>Predictive customer intelligence</li> <li>Anomaly detection</li> </ul>

[nvidia.com/en-us/gpu-cloud/](https://nvidia.com/en-us/gpu-cloud/)

## Resources

**[Galvanizing manufacturing with Artificial Intelligence and Deep Learning](#)**

**[Make data work for you with HPE solutions and H2O Driverless AI](#)**



## HPE Apollo 6500 Gen10 system: Optimized for AI and ML

Providing superior performance-per-dollar for GPU intensive workloads with up to eight high performance GPUs per server delivering up to 125 TFLOPS single precision compute. HPE takes ML workloads to new heights of performance, efficiency, and security by offering bundles that allow for customization. By applying this unique configuration for ML, you benefit from:

- HPE's industry-leading server technology that delivers unprecedented performance with reliability, availability, and services (RAS) features
- Selecting a complete AI/ML solution that delivers better price performance
- Worldwide services (Advisory, Professional, and HPE GreenLake consumption model) and support
- 3-year parts, labor, and on-site support warranty
- Simple system management
- Essential firmware is anchored by the HPE iLO 5 chip by HPE 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
- Superior performance per dollar, ease of serviceability, and zoned cooling with GPUs
- Second-generation Intel® Xeon® Scalable processors for highest levels of performance and reliability




Figure 1. HPE Apollo 6500 Gen10 with HPE ProLiant XL270d Gen10 server

- Broad choice of tier-one OS supported
- Added security and smart remote functionality with the HPE iLO 5 Advanced License
- Customization of compute, storage, and memory to fit your unique needs
- A sophisticated set of metering financial and services components delivered to you at the local level with HPE GreenLake

### Validated solution elements for training with HPE Apollo 6500

- **Servers:** (1) HPE ProLiant XL270d Gen10 server
- **CPU:** (2) Intel Xeon Scalable Gold 6148 processors
- **Memory:** (12) HPE 32GB 2Rx4 PC4-2666V-R Smart Memory and (12) HPE 16GB 2Rx8 PC4-2666V-R Smart Memory
- **Storage:** (4) HPE 1.92TB SATA MU SFF SC DS SSD
- **Smart storage controller:** (1) HPE Smart Array P408i-a Controller
- **NICs:** (1) HPE Ethernet 10/25Gb 2-port 640FLR-SFP28 Adapter and (4) HPE IB EDR/EN 100Gb 1-port 840QSFP28
- **GPUs:** (8) NVIDIA Tesla V100 GPU with 300 GB/s NVLink
- **Service and support:** 3/3/3—Server Warranty, worldwide HPE Pointnext Services (extensive Advisory, Professional, Operational, HPE GreenLake flexible consumption model), HPE iLO Advanced License

### Designed to scale

The chassis with up to eight GPUs helps you meet the demanding data read/write requirements on the storage and data management components of HPC analytics and AI environments is built to scale. Also, to accommodate your unique requirements, bundles provide the flexibility to configure the CPU, storage, and memory capacity as needed. As your needs grow, you can easily scale these solutions with additional flash storage and HPE Apollo Server nodes to support your increasing needs.

### A powerful solution supported by global services

HPE Pointnext Services leverages our strength in infrastructure, partner ecosystems, and end-to-end lifecycle experience to accelerate powerful, scalable IT solutions to provide you the assistance for faster time to value. HPE Pointnext Services provides a comprehensive portfolio including Advisory and Transformational, Professional, and Operational Services to help accelerate your digital transformation.

### Let's make it happen together

The HPE Apollo bundles for ML make it easy to configure, deploy, and manage. Get started today by sending questions or inquiries to [ai\\_madeeasy@hpe.com](mailto:ai_madeeasy@hpe.com) or contact your local HPC/AI specialist to take advantage of these bundles at special pricing.

Learn more at  
[hpe.com/servers/apollo6500](https://hpe.com/servers/apollo6500)  
[hpe.com/us/en/solutions/artificial-intelligence.html](https://hpe.com/us/en/solutions/artificial-intelligence.html)

Make the right purchase decision. Click here to chat with our presales specialists.

Share now

Get updates

© Copyright 2019 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Intel Xeon is a trademark of Intel Corporation in the U.S. and other countries. NVIDIA is a trademark and/or registered trademark of NVIDIA Corporation in the U.S. and other countries. All other third-party marks are property of their respective owners.

a50000076ENW, August 2019, Rev. 1