


HPE ProLiant Compute NVIDIA RTX PRO 4500 Blackwell Server Edition



Edge AI
& Smart spaces



Small model
inference



Video analytics



Data processing
& analytics



Broadcast &
live stream



Virtual
workstation



up to:

2X

faster AI video analysis*

3X

faster visualization
& data processing*

9.4X

faster reasoning
capabilities*

* Compared to L4

Optimized for the edge

What's this all about? NVIDIA RTX Pro is a graphics processing unit (GPU) lineup for visualization and AI inference, certified for HPE ProLiant Compute servers.

- The new NVIDIA RTX PRO™ 4500 Blackwell Server Edition GPU gives customers faster AI and graphics performance — *especially* in edge environments where power and cooling are limited.
- It helps customers speed up inference, video analytics, and real-time decision-making without needing large data-center GPUs.
- Paired with HPE ProLiant Compute, these GPUs deliver reliable, scalable AI at the edge.

What's driving the accelerated edge?

- Real-time AI + IoT growth: Technology at the edge requiring near instant, on-site AI decisions — driving demand for low-latency, local processing at the data source.
- Privacy + bandwidth efficiency: Organizations must keep certain data local to protect sensitive information, cut network costs, and meet compliance requirements.

Why HPE ProLiant Compute + NVIDIA RTX PRO 4500 Blackwell?

Sample 60 Second Pitch:

The HPE ProLiant DL145 Gen11 server now supports the new NVIDIA RTX Pro 4500 Blackwell, making it an ideal platform for GPU acceleration at the edge for AI, data processing, and virtualization.

Key strengths of HPE ProLiant DL145 Gen11:

1. Compact 2U design & short depth make it ideal for space-constrained edge environments
2. Ruggedized and equipped to handle dust and vibration
3. HPE iLO and Compute Ops Management, deliver secure, streamlined lifecycle operations across distributed sites.

If you're exploring ways to run AI or other GPU-powered workloads at the edge, the HPE ProLiant DL145 paired with the NVIDIA RTX PRO 4500 Blackwell is an efficient and highly capable system built for real-world environments.

Unlock the edge with the NVIDIA Blueprint for Agentic AI



- **Architecture:** NVIDIA agentic AI software stack with HPE ProLiant edge platforms and NVIDIA RTX PRO 4500 Blackwell.
- **Agentic AI blueprints** so agents that can reason, plan, and act rather than simply respond to prompts.

NVIDIA AI-Powered Retail Shopping Assistant



- **Personalized, multimodal shopping** with natural language and image-based product discovery.
- **Rapid deployment** using a production-ready NVIDIA blueprint that reduces design and integration effort, certified for HPE ProLiant Compute DL380 Gen12.

HPE ProLiant Compute NVIDIA RTX PRO 4500 Blackwell Server Edition

PRO TIP: Target Older Generation GPUs for Refresh

Focus on the benefits of moving from out-of-date competitor platforms running older NVIDIA edge GPUs (T4, L4, A10) or CPU-only edge systems to new, accelerated HPE ProLiant Compute platforms:

Accelerate the edge

- High performance per watt in a 165W edge form factor, helping move inference from cloud to edge and reduce operating costs.
- Faster visualization and AI video processing, improving remote productivity, Vision AI, and real-time decision-making.

Secure the edge

- HPE iLO 7 silicon root of trust helps safeguard vulnerable edge deployments against firmware and supply-chain-level attacks.

Automate the edge

- With HPE Compute Ops Management, use one cloud-based console to monitor, manage, and automate edge servers while spending up to 75% less time on server management.¹
- Lower TCO than multi-GPU or cloud-only approaches by keeping compute local and reducing bandwidth and cloud costs.

RTX Pro Anywhere: A Unified Platform for Data Center and at the Edge

Together, the NVIDIA RTX PRO 4500 Blackwell and RTX PRO 6000 Blackwell deliver a continuum of accelerated computing.

The RTX PRO 4500 Blackwell brings fast, efficient AI inference and advanced graphics with 32GB GDDR7 and next-gen Tensor/RT/CUDA cores. The RTX PRO 6000 Blackwell scales up to large generative AI, simulation, and multi-workload data-center performance with 96GB GDDR7, massive core counts, and an enterprise-optimized Blackwell architecture.

This lets organizations run everything from real-time edge AI, visualization, and design to mission-critical, memory-intensive AI model development and high-performance rendering—scaling seamlessly from workstation to data center with consistent acceleration, reliability, and efficiency.

NVIDIA RTX PRO 6000 600W 96GB 2S FHFL		BEST Performance Universal Compute + Graphics		NVIDIA RTX PRO 4500 165W 32GB 1S FHFL		Developer / Edge Universal Compute + Graphics	
							
							
HPE ProLiant Compute Secure from chip to cloud. Automated by AI. At the Edge or in the Data Center							

¹ Findings based on comparing COM to legacy/manual management approaches, these are not direct comparisons to HPE OneView



HPE ProLiant Compute


NVIDIA RTX PRO 4500 Blackwell Server Edition

	RTX PRO 6000 Blackwell Server Edition	RTX PRO 4500 Blackwell Server Edition
GPU Architecture	Blackwell RTX	Blackwell RTX
FP8 FP4 Tensor core	Yes Yes	Yes Yes
Memory	96 GB	32 GB
Memory Bandwidth	1.6 TB/s	800 GB/s
MIG Capability	Up to 4 MIG	Up to 2 MIG
Virtual Machines	Up to 48 VMs	Up to 16 VMs
Media Engines	4 Video Encoder 4 Video Decoder	3 Video Encoder 3 Video Decoder
Confidential Compute	Yes	Yes
Interconnect	PCIe Gen5 x16	PCIe Gen5 x16
Form Factor	2 Slot FHFL	1 Slot FHFL
Power	600W	165W
OCA Smart Template or UCID	DL380a: search "Enterprise AI" for small, med, large	DL145: 5152087685-01 (pre-release)
		Additional platforms soon

Objections & Responses:

- **“Cloud GPUs are cheaper.”**
Edge inference eliminates egress fees, reduces latency, and cuts ongoing cloud costs.
- **“What makes the RTX 4500 different from the RTX 6000?”**
The RTX 4500 is smaller and requires less power (only 165W) which makes it ideal for AI inference & video analytics at the edge.
- **“We already have T4/L4/A10.”**
NVIDIA’s new Blackwell architecture delivers **2–9×** performance uplift in the same power envelope.
- **“Do we even need a GPU?”**
CPUs can’t support real-time inference or multi-stream analytics at edge power budgets.

Sales preparation journey from 0 to first customer conversation



Visit [Sales Briefcase](#)

© Copyright 2026 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.

This document contains confidential and/or legally privileged information. It is intended for Hewlett Packard Enterprise and Channel Partner Internal Use only. If you are not an intended recipient as identified on the front cover of this document, you are strictly prohibited from reviewing, redistributing, disseminating, or in any other way using or relying on the contents of this document.

All third-party marks are property of their respective owners.

a00157716enw