



HPE ProLiant DL385 Gen11 with NVIDIA RTX PRO 6000 Blackwell Server Edition

Flexible data center platform to accelerate Enterprise Workloads and AI

At a glance

The HPE ProLiant DL385 Gen11 is powered by up to two AMD EPYC™ processors and up to two NVIDIA RTX PRO™ 6000 Blackwell Server Edition GPUs with a unified infrastructure to enable IT to streamline operations and accelerate time to value across generative design, LLM inference, digital twins, VDI, and immersive visualization. With support for NVIDIA® AI Enterprise, the HPE ProLiant DL385 Gen11 empowers organizations to accelerate workload deployment—securely, efficiently, and at speed.

As AI continues to redefine enterprise innovation, organizations are rapidly adopting advanced workloads spanning agentic AI, physical AI, and scientific computing—from robotics and simulation to rendering, 3D graphics, and video generation. These diverse demands require a universal data center platform capable of consolidating AI and visual computing workloads with uncompromising performance, scalability, and security.

AI and visual computing workloads

The **HPE ProLiant DL385 Gen11**, combined with up to two **NVIDIA RTX PRO 6000 Blackwell Server Edition GPUs**, delivers measurable business value by enabling enterprises to consolidate AI and visual computing workloads onto a single, universal data center platform. This approach reduces infrastructure complexity, optimizes resource utilization, and accelerates time to value for AI-driven initiatives—all while maintaining enterprise-grade security and operational efficiency.

This flexible platform supports a broad spectrum of use cases across industries:

- **Healthcare:** Accelerate medical imaging, genomics analysis, and AI-driven diagnostics
- **Manufacturing:** Power robotics, digital twins, and predictive maintenance for smart factories
- **Media and entertainment:** Enable real-time rendering, 3D graphics, and video production for immersive experiences
- **Financial services:** Enhance fraud detection, risk modeling, and algorithmic trading with AI-driven insights
- **Scientific research:** Drive breakthroughs in simulation, computational modeling, and large-scale data analysis

By supporting **agentic AI, physical AI, scientific computing, rendering, 3D VDI graphics, and video**, the DL385 with up to two NVIDIA RTX PRO 6000 GPUs provides the agility and performance needed to meet evolving enterprise demands—securely and at scale—up to:

- 45x better performance versus traditional CPU-based workloads¹
- 6x faster inference performance compared with the previous-generation NVIDIA L40S GPU¹
- 4x higher photorealistic rendering and visualization performance than the L40S GPU¹

¹ "NVIDIA RTX PRO Servers With Blackwell Coming to World's Most Popular Enterprise Systems," NVIDIA press release, 2025



NVIDIA RTX PRO 6000 Blackwell Server Edition

Built on the groundbreaking NVIDIA Blackwell architecture, the NVIDIA RTX PRO 6000 Blackwell Server Edition delivers a powerful combination of AI and visual computing capabilities to accelerate enterprise data center workloads.

Technical highlights

- **Compute power:** Fifth-generation Tensor Cores with FP4 precision deliver up to 3x AI performance over the previous generation
- **Rendering and graphics:** Fourth-generation ray tracing cores enable 2x faster rendering for photorealistic visualization and 3D design
- **Memory bandwidth:** 96 GB GDDR7 with 1.6 TB/s bandwidth supports large-scale LLM inference (up to 70B parameters) and complex simulations.
- **Multi-instance GPU (MIG):** Up to four isolated GPU instances, each with dedicated memory and compute resources
- **NVIDIA vGPU:** Up to 12 VMs per instance / 48 VMs per GPU can be shared on a single GPU
- **Video acceleration:** Ninth-gen NVENC and sixth-gen NVDEC engines for high-speed AV1 and HEVC encoding/decoding

Performance improvements enable enterprises to accelerate time to insight, reduce infrastructure sprawl, and deliver richer user experiences.

HPE ProLiant DL385 Gen11

The HPE ProLiant DL385 Gen11 is a versatile 2U server platform designed to meet the evolving demands of enterprise AI and data-intensive workloads. Built on the latest AMD EPYC processors and engineered for GPU acceleration, it provides the performance, flexibility, and security required to consolidate AI, analytics, and visual computing workloads in a single, efficient infrastructure.



Technical highlights

- **Compute power:** Supports up to two 4th or 5th Gen AMD EPYC processors for exceptional core density and performance per watt
- **GPU acceleration:** Accommodates up to two NVIDIA RTX PRO 6000 Blackwell Server Edition GPUs for advanced AI and graphics workloads
- **Scalable architecture:** Flexible configuration options for memory, storage, and networking to adapt to diverse enterprise needs
- **Optional direct liquid cooling:** For advanced energy efficiency
- **Security by design:**
 - **HPE iLO 6 with silicon root of trust:** Establishes a hardware-anchored chain of trust from power-on to runtime, protecting against firmware attacks and unauthorized access
 - **Quantum-resistant firmware signing:** Future-proof security for evolving cyber threats
- **Operational efficiency:**
 - **HPE Compute Ops Management:** Cloud-native, AI-driven automation for lifecycle management, predictive analytics, and global fleet visibility
 - **HPE OneView:** Simplifies orchestration and integration across hybrid environments

The DL385 Gen11 combines **performance, security, and manageability** to deliver a universal platform for AI and enterprise workloads. With **silicon root of trust of HPE iLO** helping ensure uncompromised security and **HPE Compute Ops Management** enabling predictive, automated operations, organizations can confidently scale innovation while reducing risk and operational complexity.

Visit [HPE.com](https://www.hpe.com)

Learn more at

[HPE ProLiant DL385 Gen11](#)

[NVIDIA RTX PRO 6000 Blackwell Server Edition](#)

HPE



[Chat now](#)

© Copyright 2025 Hewlett Packard Enterprise Development LP. The information contained herein is subject to change without notice. The only warranties 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.

AMD is a trademark of Advanced Micro Devices, Inc. NVIDIA RTX, vGPU, the NVIDIA logo, and NVIDIA are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. All third-party marks are property of their respective owners.

a00150700ENW

HEWLETT PACKARD ENTERPRISE

[hpe.com](https://www.hpe.com)