

HPE ProLiant Compute DL384 Gen12



What's new

- The first HPE ProLiant rack mount server with the NVIDIA GH200 Grace Hopper™ Superchip.
- Support for dual NVIDIA GH200 Grace Hopper Superchips with NVIDIA GH200 NVL2.
- The GH200 NVL2 fuses two Grace CPUs and two Hopper GPUs in an innovative architecture to deliver 8 petaflops of AI performance into a single node.
- The GH200 NVL2 offers up to 3.5X more GPU memory capacity and 3X more bandwidth than the NVIDIA H100 Tensor Core GPU in a single server for compute- and memory-intensive workloads. [1]
- Train, tune, and deploy generative AI models with confidence.

Overview

Does your enterprise need a server to dramatically enhance performance for AI inferencing workloads in the era of accelerated computing and generative AI with the modern NVIDIA GH200 Grace Hopper™ Superchip?

The HPE ProLiant Compute DL384 Gen12 provides outstanding coherent memory and memory bandwidth with up to two superchips per server with NVIDIA GH200 NVL2, allowing teams to run larger models faster and more cost-effectively. By partnering with NVIDIA, Hewlett Packard Enterprise continues to innovate with impressive infrastructure enhancements to help customers unlock scaled-out and accelerated computing for generative AI. With superchip performance in your AI factory, HPE ProLiant Compute DL384 Gen12 delivers superb performance per GPU in the HPE ProLiant portfolio. Simplify your approach to scaling with data demands and integrate accelerated compute to reliably prepare your data center for mixed workloads and retrieval-augmented generation (RAG) in the modern hybrid reality.

Features

Modernize Your Infrastructure for the Future of AI

The HPE ProLiant Compute DL384 Gen12 features NVIDIA GH200 NVL2 Grace Hopper Superchips servers that dramatically increase memory space and bandwidth when compared to a single superchip server.

With up to with 1.2 TB of combined fast memory, a high-bandwidth connection between the NVIDIA Grace CPU and Hopper GPU can help scientists and researchers uncover unprecedented solutions to complex problems by easily handling larger LLMs and supporting applications that run terabytes of data.

Achieve greater performance per dollar invested on inference by connecting two GH200 superchips with NVLink™ in a single node to boost innovation, increase resource utilization, leverage data to enhance efficiency, and improve ROI.

Innovative architecture allows lower latency of responses and enables computationally-demanding workloads to easily run on fewer nodes, reducing the costs associated with maintaining complex distributed architectures and multiple operating systems.

Unlock More Value and Accelerate AI Outcomes for Better Insight and Innovation

With support for the modern NVIDIA InfiniBand, Ethernet, and BlueField adapters, the HPE ProLiant Compute DL384 Gen12 offers reliable, low latency and high-speed interconnects for faster inferencing.

Customize performance and accelerate workloads from the data center to the cloud by utilizing HPE ProLiant servers which are specifically engineered for the hybrid world.

As more and more processes become AI-supported and data-driven, leverage an AI-native architecture to capitalize on the AI revolution and enable customers to easily deploy AI into their own business processes.

Flexible Management Keeps You in Control

Building on HPE ProLiant as the legendary foundation, the HPE ProLiant Compute DL384 Gen12 delivers a consistent experience across the HPE ProLiant portfolio with an HPE iLO management and firmware stack to facilitate robust reliability and security with Silicon Root of Trust technology from HPE.

The included HPE Integrated Lights-Out 6 (iLO 6) server management software enables enhanced security when configuring, monitoring, and updating your HPE ProLiant Gen12 server seamlessly from anywhere in the world.



Technical specifications		HPE ProLiant Compute DL384 Gen12
Processor family	NVIDIA Grace CPU and Hopper GPU	
Power supply type	Up to 4 HPE 1800W-2200W Flex Slot Titanium Hot Plug Power Supply Kits For a single GH200: 2+0, 2+1, or 2+2 redundancy. Minimum of two power supplies. For a dual GH200 with NVL2: 3+1 redundancy, and 4 power supplies required, no other choices supported.	
Expansion slots	Up to four high speed PCIe Gen5 x16 devices, Up to four PCIe Full Height Half Length (FHHL) slots, or Up to two PCIe FHHL and two OCP 3.0	
System fan features	6 hot plug redundant fans	
Storage controller	HW RAID controllers are not supported this generation for either the M.2 or EDSFF NVMe drives	
Infrastructure management	HPE iLO Standard with intelligent provisioning (embedded), HPE iLO Advanced	
Warranty	3/3/3: Server warranty includes three years of parts, three years of labor, and three years of onsite support coverage. Additional information regarding worldwide limited warranty and technical support is available at: https://h20564.www2.hpe.com/hpsc/wc/public/home . Additional HPE support and service coverage, to supplement the product warranty, is available. For more information, visit https://www.hpe.com/support .	
Drive supported	Up to 8 EDSFF NVMe Gen5 drives	

[1] Visit URL [nvidia.com/en-us/data-center/grace-hopper-superchip](https://www.nvidia.com/en-us/data-center/grace-hopper-superchip)



For additional technical information, available models and options, please reference the QuickSpecs

HPE Services

No matter where you are in your transformation journey, you can count on HPE Services to deliver the expertise you need when, where and how you need it. From strategy and planning to deployment, ongoing operations and beyond, our experts can help you realize your digital ambitions.

Consulting services

Experts can help you map out your path to hybrid cloud and optimize your operations.

Managed services

HPE runs your IT operations, giving you unified control, so you can focus on innovation.

Operational services

Optimize your entire IT environment and drive innovation. Manage day-to-day IT operational tasks while freeing up valuable time and resources.

- HPE Complete Care Service: a modular service designed to help optimize your entire IT environment and achieve agreed upon IT outcomes and business goals. All delivered by an assigned team of HPE experts.
- HPE Tech Care Service: the operational service experience for HPE products. The service provides access to product specific experts, an AI driven digital experience, and general technical guidance to help reduce risk and search for ways to do things better.

Lifecycle Services

Address your specific IT deployment project needs with tailored project management and deployment services.

HPE Education Services

Training and certification designed for IT and business professionals across all industries. Create learning paths to expand proficiency in a specific subject. Schedule training in a way that works best for your business with flexible continuous learning options.

The Defective Media Retention (DMR) service feature option applies only to Disk or eligible SSD/Flash Drives replaced by Hewlett Packard Enterprise due to malfunction. Comprehensive Defective Material Retention (CDMR) allows you to keep all data retentive components.

HPE GreenLake

HPE GreenLake edge-to-cloud platform is HPE's market-leading as-a-Service offering that brings the cloud experience to apps and data everywhere – data centers, multi-clouds, and edges – with one unified operating model, on premises, fully managed in a pay per use model.

If you are looking for more services, like **IT financing solutions**, please explore them [here](#).

Make the right purchase decision.
Contact our presales specialists.

[Find a partner](#)



Share now



Get updates

Explore **HPE GreenLake**



© Copyright 2024 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.

Parts and Materials: HPE will provide HPE-supported replacement parts and materials required to maintain the covered hardware.

Parts and components that have reached their maximum supported lifetime and/or the maximum usage limitations as set forth in the manufacturer's operating manual, product quick-specs, or the technical product data sheet will not be provided, repaired, or replaced as part of these services.

Arm is a registered trademark of Arm Limited. Ubuntu is a registered trademark of Canonical Ltd. NVLink 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.

Image may differ from the actual product
[PSN1014835194WWEN](#), November, 2024.