

NVIDIA Accelerators for HPE

NVIDIA L4 24GB PCIe Accelerator for HPE (S0K89C)



What's new

- NVIDIA H200 NVL is ideal for air-cooled enterprise rack designs that require flexible configurations.
- With up to four GPUs connected by NVIDIA NVLink™ and a 1.5x memory increase, LLM inference can be accelerated up to 1.7x and HPC up to 1.3x on H200 NVL compared to H100 NVL.
- NVIDIA H200 NVL comes with a five-year NVIDIA AI Enterprise subscription and simplifies the way you build an enterprise AI-ready platform.
- Third generation NVLink doubles the GPU-GPU direct bandwidth.
- Third-generation RT cores for speeding up rendering workloads.

Overview

Do you require higher performance for artificial intelligence (AI) training and inference, high-performance computing (HPC) or graphics? NVIDIA® Accelerators for HPE help solve the world's most important scientific, industrial, and business challenges with AI and HPC. Visualize complex content to create cutting-edge products, tell immersive stories, and reimagine cities of the future. Extract new insights from massive datasets. Hewlett Packard Enterprise servers with NVIDIA accelerators are designed for the age of elastic computing, providing unmatched acceleration at every scale.

Features

Develop and Deploy AI at Any Scale

Build new AI models with supervised or unsupervised training for generative AI, computer vision, large language models (LLM), scientific discovery, and financial market modeling with NVIDIA accelerators and HPE Cray systems.

Get real-time inference for computer vision, natural language processing, fraud detection, predictive maintenance, and medical imaging with NVIDIA

- NVIDIA confidential computing is a built-in security feature of Hopper-based accelerators.

accelerators and HPE ProLiant Compute Servers.

NVIDIA Accelerators for HPE improve computational performance, dramatically reducing the completion time for parallel tasks, offering quicker time to solutions.

NVIDIA Qualified and NVIDIA Certified HPE Servers

NVIDIA Accelerators for HPE undergo thermal, mechanical, power, and signal integrity qualification to validate that the accelerator is fully functional in the server. NVIDIA Qualified configurations are supported for production use.

NVIDIA-Certified HPE servers are tested to validate both multi-GPU and multi-node performance for a diverse set of workloads to deliver excellent application performance, manageability, security, and scalability.

HPE Integrated Lights-Out (iLO) Management

HPE iLO server management software enables you to securely configure, monitor, and update your NVIDIA Accelerators for HPE seamlessly, from anywhere in the world.

HPE iLO is an embedded technology that helps simplify server and accelerator set up, health monitoring, power, and thermal control, utilizing HPE's Silicon Root of Trust.

Technical specifications

NVIDIA L4 24GB PCIe Accelerator for HPE

Product Number	S0K89C
Platform supported	HPE Mainstream Compute Platforms
Number of accelerators per card	1
Memory size per board	24 GB GDDR6
Memory bandwidth for board	300 GB/s
Accelerator applications	AI, video, virtual workstations, and graphics applications in the enterprise, in the cloud, and at the edge.
Architecture features	The NVIDIA L4 Tensor Core GPU delivers a versatile platform to accelerate Deep Learning, Graphics and Video processing applications in the Cloud and at the Edge. It is a half-height (Low profile), Half-length, single slot card featuring 24GB of GDDR6 memory, x16 PCIe Gen4 connectivity at a 72W maximum power envelope that is passively cooled. NVIDIA L4 provides revolutionary multi-precision performance to accelerate deep learning and machine learning training and inference, video transcoding, AI audio and video effects, rendering, data analytics, virtual workstations, virtual desktop, and other workloads.
System	Compatible with HPE DL320 Gen11
Product dimensions (metric)	16.85 x 6.89 x 1.87 cm
Weight	0.27 kg
Warranty	For details on HPE Qualified Options Limited Warranty visit: 3-year parts, 0-year labor, and 0-year on-site support coverage. For more warranty information refer to https://h20564.www2.hpe.com/hpsc/wc/public/home



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

Make the right purchase decision.
Contact our presales specialists.

[Find a partner](#)



Chat now (sales)



Call now



Buy now



Share now



Get updates


**Hewlett Packard
Enterprise**

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.

Advisory & Professional 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.

Support 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.
- **HPE Multivendor Services:** Single point of accountability for managing on-site hardware and software support for multivendor products. HPE experts help manage your IT across technologies and platforms for HPE and non-HPE technologies, acting as the single point of contact for your IT operational needs.

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

Defective Media Retention is optional and allows you to retain Disk or eligible SSD/Flash Drives replaced by HPE due to malfunction.

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](#).

Visit **HPE.com**



© Copyright 2025 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.

NVIDIA, the NVIDIA logo, NVIDIA-Certified Systems, NVLink, NVIDIA Omniverse and NVIDIA RTX are trademarks and/or registered trademarks of NVIDIA Corporation and affiliates in the U.S. and other countries. Other company and product names may be trademarks of the respective owners with which they are associated.

Image may differ from the actual product
[PSN1014731952CZEN](#), April, 2025.