

HPE ProLiant Compute XD680

HPE ProLiant Compute XD680 Air Cooling Server (S3W15A)



What's new

- Eight Intel® Gaudi® 3 AI accelerators that deliver strong performance for AI training, tuning and inferencing.
- 5U form factor, air-cooled.
- 24x DDR5 RDIMMs with 128 GB HBM capacity and 3.7 TB/sec HBM bandwidth.
- HPE iLO 6 offers seamless management and built-in security features.
- Support for 5th Gen Intel® Xeon® Scalable processors, up to 400W TDP.

Overview

Do you need a large, high-performant and secure AI cluster? The HPE ProLiant Compute XD680 is designed to address demanding AI training, tuning, and inferencing workloads. It leverages a cost-effective, purpose-built architecture and features eight Intel® Gaudi® 3 AI accelerators, along with a portfolio of advanced software and services. This server is customized for GPU-intensive workloads, such as large language model (LLM) training and natural language processing (NLP), to gain insights from large amounts of data. Managed by HPE iLO, it offers seamless management and robust security. HPE Services provides expertise for global installation and deployment of large AI clusters, enabling stability and operational excellence.

Features

End-to-End Customized Solution to Deploy Cost-Effective and Protected Large AI Clusters

HPE ProLiant Compute XD680 features eight Intel® Gaudi® 3 AI accelerators that deliver customized price-performance and the ability to efficiently scale from 1 to 100s of nodes to meet your generative (Gen AI) and LLM requirements.

HPE iLO offers seamless server management with built-in security innovation to protect against escalating security threats.

HPE Performance Cluster Manager expedites initial cluster setup and software deployment, and provides continued resilience and reliability.

HPE Cray Storage Systems C500 provides very strong read and write performance without capacity-based licensing, removing any I/O bottlenecks holding your AI initiatives back and enabling you to take control of your AI storage spending.

Global Solution and Expertise

The robust and secure global supply chain robust and secure global supply chain from Hewlett Packard Enterprise, and our proven ability to deploy massive custom AI clusters anywhere in the world with superior serviceability enables you to establish your presence globally with confidence.

You get flexibility and agility with a customized IT environment that lets you run the right workload at the right place, at the right time. Whether you want a solution on-premises, or deploy a hybrid model, we offer flexible consumption and deployment models to suit your requirements and goals.

The HPE Center of Excellence for AI and data helps you define and refine your AI business opportunities with advice from a team of experts located around the world.

HPE Tech Care Service combines deep technical expertise with a leading-edge digital experience to provide personalized, prompt, and reliable product support, so you can get the most out of your global IT investments.

Speed and Efficiency to Increase your Competitive Advantage

Comprehensive portfolio of HPE Services to build, test, and deploy your solution with ongoing support, all available as stand-alone or consumption-based services, accelerates the deployment of your environment.

HPE Factory Express Services builds, integrates, validates, fully tests, and customizes your solution in the factory, facilitating quicker on-site deployment.

HPE Professional and Advisory services, offering best-in-class methodologies and tools to enhance performance and increase ROI, help you navigate the challenges of a large AI deployment.

HPE Financial Services allows you to turn aging infrastructure into capital in a secure, globally consistent and environmentally responsible way, without disrupting your AI production.



Technical specifications

HPE ProLiant Compute XD680 Air Cooling Server

Product Number	S3W15A
Processor number	<p>5th Generation Intel® Xeon® Scalable processors</p> <p>Specific SKUs supported:</p> <p>Intel Xeon-Gold 6548Y+ 2.5GHz 32-core 250W Processor for HPE</p> <p>Intel Xeon-Gold 6544Y 3.6GHz 16-core 270W Processor for HPE</p> <p>Intel Xeon-Platinum 8558P 2.1GHz 48-core 330W Processor for HPE</p> <p>Intel Xeon-Platinum 8562Y+ 2.8GHz 32-core 300W Processor for HPE</p> <p>Intel Xeon-Platinum 8568Y+ 2.3GHz 48-core 350W Processor for HPE</p> <p>Intel Xeon-Platinum 8570 2.1GHz 56-core 350W Processor for HPE</p> <p>Intel Xeon-Platinum 8580 2.0GHz 60-core 350W Processor for HPE</p> <p>Intel Xeon-Platinum 8592+ 1.9GHz 64-core 350W Processor for HPE</p>
Processor core available	16–64 cores per CPU
Processor cache	<p>L1 Cache: 32KB (instruction), 48KB (data) per core</p> <p>L2 Cache: 1.25MB per core</p> <p>L3:Cache: Varies per below:</p> <p>Intel Xeon-Gold 6548Y+: 60MB / core</p> <p>Intel Xeon-Gold 6544Y: 45MB / core</p> <p>Intel Xeon-Platinum 8558P: 260MB / core</p> <p>Intel Xeon-Platinum 8562Y+: 60MB / core</p> <p>Intel Xeon-Platinum 8568Y+: 300MB / core</p> <p>Intel Xeon-Platinum 8570: 300MB / core</p> <p>Intel Xeon-Platinum 8580: 300MB / core</p> <p>Intel Xeon-Platinum 8592+: 320MB / core</p>
Processor speed	1.9 GHz–3.6 GHz
Power supply type	<p>3000W CRPS Titanium, 200-277 VAC input</p> <p>54V output</p>
Expansion slots	8x HHHL x16 PCIe Gen5 slots, 4x FHHL x16 PCIe Gen5 slots, 2x OCP3 slots
Persistent memory type	N/A
Persistent memory media capacity	N/A
Memory type	DDR5 RDIMMs
Optical drive type	N/A
Network controller	<p>The HPE ProLiant Compute XD680 contains a 1GbE LAN for HPE iLO control through a standard RJ-45 connector</p> <p>The following NICs are supported:</p> <p>BCM 57504 10/25GbE 4p SFP28 OCP3 Adptr</p> <p>INT E810 10/25GbE 2p SFP28 Adptr</p> <p>INT E810 10/25GbE 2p SFP28 OCP3 Adptr</p> <p>BCM 57416 10GbE 2p BASE-T OCP3 Adptr</p> <p>INT E810 100GbE 2p QSFP28 Adptr</p> <p>INT I350 1GbE 4p BASE-T OCP3 Adptr</p> <p>INT I350 1GbE 4p BASE-T Adptr</p> <p>INT E810 100GbE 2p QSFP28 OCP3 Adptr</p> <p>MLX MCX623105AS 200GbE 1p QSFP56 Adptr</p>



Storage controller

Boot: HPE NS204i-u Gen11 NVMe Hot Plug Boot Optimized Storage Device ("Wondermiddel")

HPE MR416i-o Gen11 SPDM Storage Cntlr
HPE MR416i-p Gen11 SPDM Storage Cntlr
HPE SR932i-p Gen11 SPDM Storage Cntlr

Infrastructure management

HPE iLO Advanced
HPE iLO Advanced licenses offer smart remote functionality without compromise, for all HPE ProLiant servers. The license includes the full integrated remote console; virtual keyboard, video, and mouse (KVM); multiuser collaboration; console record and replay; and GUI-based and scripted virtual media and virtual folders. You can also activate the enhanced security and power management functionality.

HPE Performance Cluster Manager (HPCM)
Designed to operate from dozens of nodes to > 10,000 nodes in a cluster. HPCM dashboards not only show system performance but also:
Cluster Health at-a-Glance–Single Pane view for the complete cluster Health Status
Live System Monitoring–Dashboards for key metrics like power, cooling, CPU, memory, disk, fabric, cluster, and job scheduler monitoring metrics.

Warranty

This product is covered by a global limited warranty and supported by HPE Services and a worldwide network of HPE Authorized Channel Partners resellers. Hardware diagnostic support and repair is available for three years from date of purchase. Support for software and initial setup is available for 90 days from date of purchase. Enhancements to warranty services are available through HPE Services operational services or customized service agreements. Hard drives have either a one year or three-year warranty; refer to the specific hard drive QuickSpecs for details.

Notes: Server Warranty includes 3-Year Parts, 3-Year Labor, 3-Year On-site support with next business day response.

© 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

Intel® and Xeon® are registered trademarks of Intel Corporation in the U.S. and other countries.




For hard drives, 1GB = 1 billion bytes. Actual formatted capacity is less



For additional technical information, available models and options, please reference the [QuickSpecs](#)

Make the right purchase decision.
Contact our presales specialists.

[Buy now](#)



**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.

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 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](#).

Explore **HPE GreenLake**



© 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.

Intel, Intel Xeon, Intel Xeon Gold, and Intel Xeon Platinum are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. MLX is either a registered trademark or trademark of Microsoft Corporation in the United States and/or other countries. All third-party marks are property of their respective owners.

Image may differ from the actual product
[PSN1014856853USEN](#), January, 2025.