

Accelerate large-scale AI with sustainability and security with HPE and AMD

HPE ProLiant Compute XD685



Prepare for the growing demands of AI

This is an era where artificial intelligence (AI) is a key driver for growth and innovation. AI inspires organizations worldwide to work faster and smarter than before, leveraging data to make more accurate decisions and enable better business results. As the impact of AI continues to scale, the challenges brought on by surging data volumes and the complexities of managing new technologies (with limited budget and expertise) are quickly outpacing the abilities of traditional compute environments. AI will only ever be as good as the infrastructure that supports it. For that reason, organizations need infrastructure engineered for the demands of AI. New systems are having a major impact on their success, from massive expansion in data center capacity to sustainable performance and cooling for large-scale AI — and this is only the beginning.

As power requirements for AI increase, rack density and cooling requirements have increased simultaneously. Advances in compute and chip designs cause equipment racks to double in power density every six to seven years. Greater rack density enables greater data capacity. However, it also results in higher energy use and heat generation.

AI service providers and large model builders in particular need an efficient approach to how they utilize power and cool racks. These organizations must quickly set up large AI clusters to serve the needs of their growing customer base and to stay competitive in

a dynamic market. They depend on integrated, stable, and secure compute environments that can be deployed quickly and accelerate AI performance on a massive scale.

Discover the next generation of AI systems

Hewlett Packard Enterprise and AMD understand the challenges of AI. We know that training large language models (LLMs) requires strong scaling and massively parallel compute capabilities. As leaders in delivering the [world's most powerful and energy-efficient systems](#), HPE and AMD are extending our innovative technologies and services to support a growing market of AI model builders, as well as accelerate scientific and engineering breakthroughs across industries. Our global presence and deep expertise are empowering AI service providers and ambitious enterprises building and training large AI models to deploy AI clusters anywhere in the world. We offer financial and deployment flexibility to enable these organizations to quickly build large and scalable AI environments.

We are eager to announce a powerful new solution that accelerates AI training with a purpose-built architecture and a portfolio of advanced software and services. [HPE ProLiant Compute XD685](#) is designed to meet escalating AI performance requirements and advance your sustainability goals. It is the first offering within a new family of AI systems to have an HPE-designed chassis.




Figure 1. The HPE advantage for AI training at scale

Featuring eight of the latest [AMD Instinct™ MI355X accelerators](#) and two [5th Generation AMD EPYC™ processors](#), the system is optimized for intensive AI workloads—such as natural language processing, LLM training, and multimodal training. The choice of direct liquid cooling (DLC) increases energy efficiency, lowers operational costs, and advances sustainability. The choice of air cooling is available as an alternative. For a seamless operating experience, these systems include built-in security features delivered via [HPE iLO](#) as well as support from [HPE Services](#) experts who help you deploy large AI clusters with confidence.

Secure a competitive edge for AI

HPE ProLiant Compute XD685 systems powered by AMD offer distinct advantages for large AI workloads.

Increase competitive advantage with unparalleled speed:

- **Run the largest AI models with fewer GPUs.** AMD Instinct MI355X accelerators deliver leadership AI performance with up to 4x generational uplift, massive memory capacity, and energy-efficient support.
- **Accelerate the deployment and operations.** The trusted ecosystem of HPE Services helps to build, integrate, validate, and customize AI environments at the largest scale to facilitate quicker on-site deployment and deliver value out of the box.
- **Increase large cluster performance.** HPE provides unrivaled performance tuning capabilities for the world's largest computing systems.
- **Optimize cluster deployment.** In liquid-cooled configurations, HPE ProLiant Compute XD685 features a 5U chassis and can be deployed in a compact 8-node-per-rack arrangement

Support global footprint and growth ambitions:

- **Innovate with industry-leading technologies.** Tap into the same technologies used to power [some of the world's fastest and most sustainable supercomputers](#) and make AI work for your business. HPE and AMD solutions are delivered at scale to power all your AI initiatives. with CASB.
- **Establish your presence anywhere in the world.** HPE has a proven ability to deploy massive custom AI clusters worldwide with superior serviceability, enabled by our secure global supply chain.
- **Define and refine your AI business opportunities.** HPE experts are located worldwide to help organizations do more with AI. The [HPE Center of Excellence for AI and Data](#) in Madrid is a centralized source for data-first modernization services, empowering organizations globally to apply AI in cutting-edge, ethical, secure ways.
- **Turn aging infrastructure into cash.** [HPE Accelerated Migration](#) and [HPE Asset Upcycling Services](#) create investment capacity to accelerate transformation. Convert existing, owned IT assets into an incremental cash source or recover value from any data center technologies in a secure, globally consistent, and environmentally responsible way without disrupting AI production.

Deploy sustainable and secure large AI clusters:

- **Improve efficiency with DLC.** HPE ProLiant Compute XD685 offers the choice of DLC or air cooling. Backed by HPE's market leadership in developing the [world's largest liquid-cooled IT environments](#), organizations can deploy sustainable large AI clusters to scale with the demands of AI.




Figure 2. What makes HPE ProLiant Compute XD685 an end-to-end solution

- **Protect against growing security threats.** AMD secure processor and HPE iLO embedded server management provides “trusted security by design” for every HPE ProLiant Compute XD685. New [HPE iLO 6](#) lets organizations seamlessly configure, monitor, and update servers from anywhere to defend against increasingly complex threats.
- **Simplify management of AI clusters.** Expedite cluster setup and software deployment and ensure continued resilience with [HPE Performance Cluster Manager \(HPCM\)](#). HPCM is a fully integrated system management software that gets complex systems up-and-running quickly with automated setup from bare metal and keeps clusters healthy and operational with detailed telemetry, GPU stress tests, and more.

Prepare for the future of AI

AI is already influencing data center design, and it is expected to drive a sixfold increase in data center capacity over the next three years. The race is on to discover the promise of AI and its potential to dramatically improve outcomes in workforce productivity, healthcare, climate sciences, and much more.

HPE and AMD are committed to helping you capture this potential. We continue to deliver solutions that are purpose-built to meet the demands of AI and evolve for future requirements. Whether you are an AI service provider or an ambitious enterprise building a large AI model, HPE and AMD can transform your environments with flexible, high performance solutions that can be brought to market quickly.

Let us propel your AI ambitions to new heights.

Learn more at

[HPE.com/partners/AMD](https://hpe.com/partners/AMD)

HPE

AMD

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

[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 and the AMD Arrow logo are trademarks of Advanced Micro Devices, Inc. All third-party marks are property of their respective owners.

a50011737ENW, Rev. 2

HEWLETT PACKARD ENTERPRISE

hpe.com

