

JANUARY 2026

HPE ProLiant Compute for the Edge in a Hybrid World

Tony Palmer, Practice Director

Edge Infrastructure Challenges


As organizations race to outpace the competition in digital business, the edge represents an increasingly valuable location for investment and innovation. These locations, often separated from traditional IT systems and support and residing closer to the customer, create opportunities beyond those available with other deployment locations; however, securing and connecting these environments is a constant concern.

AI is profoundly affecting nearly every facet of how businesses operate. Infrastructure is facing heightened requirements for massive data storage capacity to support AI projects while keeping the environment secure and meeting performance demands. The most common approach to deploying generative AI is hybrid (33% of organizations),¹ which makes sense because it enables organizations to leverage scalable cloud resources for demanding workloads while maintaining control over sensitive data on premises from the core to the edge.

AI workloads—and hence infrastructure—run in a variety of locations. This means AI is a hybrid workload for many organizations, spanning a range of environments, with 43% of organizations surveyed by Enterprise Strategy Group (now Omdia) identifying the edge as an infrastructure location for AI initiatives (Figure 1).²

Figure 1. Hybrid Infrastructure Is Core to AI

Which of the following infrastructure locations is your organization using, or planning to use, for its AI initiatives? (Percent of respondents, N=350, multiple responses accepted)



Source: Omdia

¹ Source: Enterprise Strategy Group (now Omdia) Research Report, *Navigating Build-versus-buy Dynamics for Enterprise-ready AI*, January 2025.

² Source: Enterprise Strategy Group (now Omdia) Research Report, *The Critical Role of Storage in Building an Enterprise AI Infrastructure*, September 2025.

Cutting costs and boosting operational efficiency have historically been the main reasons organizations turn to edge deployments. These environments also provide organizations with additional agility and control, with the goal of consistent management of applications and infrastructure. Edge computing is no longer just about saving money; it's about unlocking strategic growth, securing a competitive advantage, and providing exceptional customer experiences.

Organizations across various industries encounter a shared challenge in determining the optimal placement of workloads: what should be deployed at the edge, what should reside in the cloud, and how to effectively implement a hybrid strategy. Factors like latency, security, compliance, and cost must be carefully evaluated to make informed decisions about workload deployment. While specific requirements may vary, the fundamental considerations for edge, cloud, and hybrid deployment models remain consistent.

Deciding where to run a workload involves evaluating several key criteria. Time-critical processes requiring ultra-low latency and deterministic responses, such as robotics or telecom radio units, are best suited for edge computing, while less time-sensitive tasks can be handled in centralized cloud environments. High-volume data sources, like video streams or sensor feeds, benefit from preprocessing at the edge to reduce cloud transmission costs, with deeper analytics performed in core data centers. Limited or unreliable WAN connectivity necessitates autonomous edge operations to ensure functionality during network outages.

Regulatory requirements around data sovereignty and compliance often mandate on-premises or in-country data storage, supported by hybrid clouds for aggregation and auditing. Remote locations with environmental constraints, such as extreme temperatures or limited space, require ruggedized edge servers designed for reliability. Security and resilience are critical, with silicon-rooted security and centralized policy control protecting edge nodes from cyberthreats while ensuring continuity during cloud disruptions. Managing large-scale edge deployments demands centralized, automated lifecycle management for provisioning, updates, and incident response. Finally, economic and sustainability considerations call for a hybrid approach that optimizes workload placement to minimize costs, power consumption, and carbon footprint.

A hybrid architecture that places each workload where it best meets security, performance, and cost objectives, with edge computing handling real-time and local needs and centralized processing and aggregation handled in the cloud or on premises is the optimal solution.

HPE ProLiant Compute Edge Portfolio

To power workloads at the edge, enterprises need robust, easy-to-manage servers that can live outside the traditional data center. The HPE ProLiant edge portfolio provides a range of such servers, each optimized for different edge scenarios. HPE designed the portfolio to address a wide range of industries' edge needs, including ruggedized designs for harsh sites, embedded silicon-based security, and unified cloud-native management.

Figure 2. HPE ProLiant Compute – Edge Computing Portfolio



Source: HPE and Omdia

HPE helps organizations manage and secure thousands of distributed servers with limited IT on-site staff.

HPE Compute Ops Management is a cloud-based software service designed to enable organizations to monitor, manage, and gain visibility into their distributed compute environments. Compute Ops Management provides lifecycle management for HPE ProLiant Compute servers. Compute Ops Management delivers automated firmware updates, OS deployment, automatic case creation, and more. Using Compute Ops Management, customers can manage their HPE ProLiant servers across their entire IT ecosystem.

HPE iLO is remote server management that uses HPE-designed silicon that prevents compromised firmware from running—a critical defense since edge sites could be physically accessed by malicious actors. iLO also enables customers to securely configure, monitor, and update HPE servers. iLO 7 features authentication using the Security Protocol and Data Model (SPDM) certification, a key security capability in servers to ensure the component’s integrity.

HPE ProLiant servers are purpose-built to accommodate accelerators. For example, the DL145 Gen11 server supports up to three single-width GPUs—like the NVIDIA L4—in a short-depth chassis. The portfolio includes both Intel and AMD-based systems like the NVIDIA L4 to suit different workloads.

Multiple models are engineered for diverse environments. The EL8000 is a modular system that can be deployed in telecom cabinets, featuring high shock/vibration resistance, front-access modularity for quick swapping of server blades, dust filtering, and can run on –48V DC power. The DL145 Gen11 server is about 50% shorter in depth than a DL365 rack server and can be mounted in tight spaces while being quiet enough (~55 dB) for office environments.

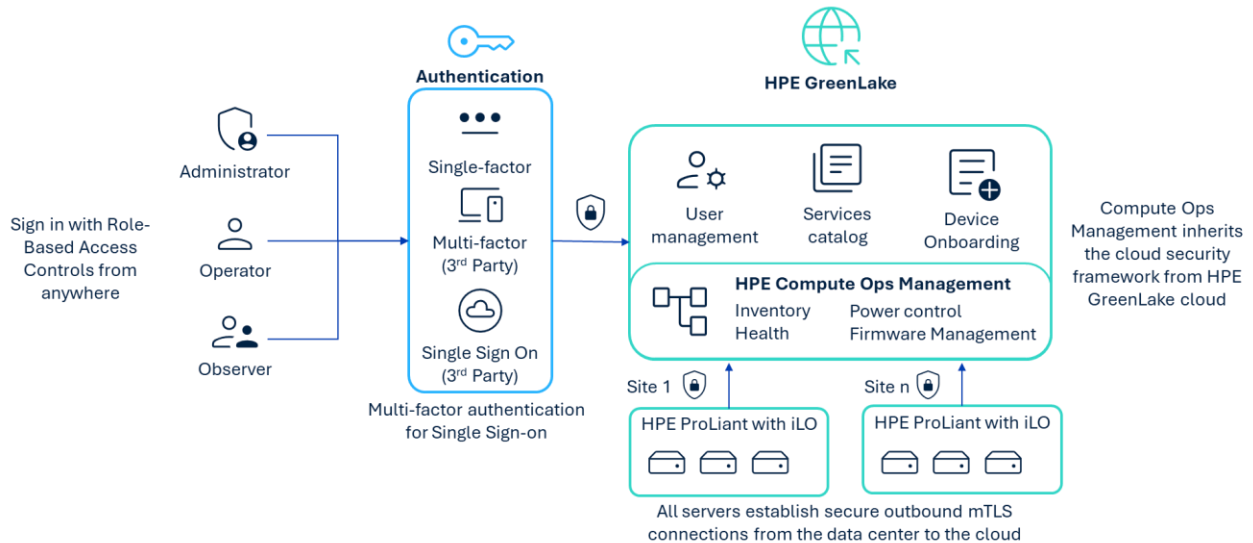
Omdia First Look

Omdia completed a technical analysis of the HPE ProLiant Compute edge computing portfolio of servers. Our goal was to validate the overall business impact of running workloads on HPE ProLiant Compute servers in hybrid and edge environments with particular focus on management, security, and AI.

Enterprises deploying industry-specific workloads at the edge require reliable, easy-to-manage servers capable of operating outside traditional data centers. The HPE ProLiant edge portfolio offers a variety of servers tailored to different edge scenarios, featuring ruggedized designs for challenging environments, embedded silicon-based security, and unified cloud-native management. These features address key considerations for deploying and managing distributed edge infrastructure. The portfolio provides flexibility to select the appropriate server for each site, such as HPE ProLiant MicroServer for small remote locations, HPE ProLiant DL20 1U racks for larger branch environments, and HPE ProLiant DL145 servers at regional hubs for AI processing. All servers in the portfolio share common management and security capabilities, running HPE iLO firmware with silicon root of trust and enabling centralized administration through the HPE Compute Ops Management platform.

We looked at how Compute Ops management facilitates device onboarding and management of inventory, health, power control, and firmware management with an inherited cloud security framework.

Figure 3. HPE Compute Ops Management



Source: HPE and Omdia

Compute Ops Management offers a range of features and benefits designed to streamline operations and enhance efficiency. It utilizes delta-based firmware updates, which identify and update only necessary components, significantly reducing download sizes and maintenance times while eliminating the need for large downloads and extended maintenance windows. Automated pre-checks ensure all parameters, such as open ports and baseline compatibility, are configured before updates. Flexible scheduling allows firmware updates to be planned up to 12 months in advance, with options for immediate updates or admin-controlled reboots.

Security management is enhanced through continuous monitoring of server fleet compliance, risk flagging, and recommendations for resolution, supporting 12 ILO security settings like password complexity, secure boot, and role-based access control. Group-based management enables centralized control of firmware baselines, power consumption, and security settings across server groups. Additionally, remote management provides console access to edge locations, reducing the need for on-site IT resources. The platform also supports seamless integration and scalability to meet the needs of growing infrastructures.

Compute Ops Management is designed with robust integration and scalability features to support diverse and large-scale infrastructures. It includes a secure gateway, a virtual appliance that aggregates ILO connections on premises, minimizing outbound connections to a single secure link, which is especially advantageous for large-scale deployments like financial institutions managing thousands of servers. The platform also offers an MSP view, enabling MSPs to monitor and manage multiple customer infrastructures through a single interface. Additionally, Compute Ops Management integrates seamlessly with tools such as HPE Aruba Networking Central, DSCC, ServiceNow, and HPE OpsRamp Software, ensuring unified management across compute, network, and storage environments.

Conclusion

AI is transforming business operations, driving the need for secure, high-performance infrastructure with massive storage capacity, and prompting organizations to adopt hybrid deployment models that balance scalable cloud resources with on-premises control of sensitive data.

Hewlett Packard Enterprise offers a broad portfolio of ruggedized, compact, and high-performance servers designed to deliver enterprise-grade security, management, and performance closer to where data is generated. HPE ProLiant servers stand out in edge environments not only for their durable physical designs but also for their embedded intelligence and security features. While certain models are built to endure harsh conditions, the consistent management and protection capabilities across the entire portfolio provide significant value for edge deployments.

Compute Ops Management simplifies server lifecycle management through automation, intelligent updates, and centralized control and addresses challenges in hybrid and edge environments, such as limited bandwidth, security compliance, and remote management. Features like Secure Gateway and MSP View make Compute Ops Management suitable for large-scale and multi-customer deployments. Compute Ops Management incorporates customer feedback to address real-world challenges, such as reducing outbound connections and providing detailed security insights.

The edge is a strategic component of modern IT architectures. Organizations shouldn't view edge and cloud as an either/or but concentrate on aligning each workload to its optimal placement and having a hybrid strategy to tie it all together. If your organization wants to deploy the right edge infrastructure to deliver real-time intelligence and services at the point of need, you should seriously consider partnering with HPE and taking a close look at the HPE ProLiant Compute edge computing portfolio.

Copyright notice and disclaimer

The Omdia research, data, and information referenced herein (the “Omdia Materials”) are the copyrighted property of TechTarget, Inc. and its subsidiaries or affiliates (together “Informa TechTarget”) or its third-party data providers and represent data, research, opinions, or viewpoints published by Informa TechTarget and are not representations of fact.

The Omdia Materials reflect information and opinions from the original publication date and not from the date of this document. The information and opinions expressed in the Omdia Materials are subject to change without notice, and Informa TechTarget does not have any duty or responsibility to update the Omdia Materials or this publication as a result.

Omdia Materials are delivered on an “as-is” and “as-available” basis. No representation or warranty, express or implied, is made as to the fairness, accuracy, completeness, or correctness of the information, opinions, and conclusions contained in Omdia Materials.

To the maximum extent permitted by law, Informa TechTarget and its affiliates, officers, directors, employees, agents, and third-party data providers disclaim any liability (including, without limitation, any liability arising from fault or negligence) as to the accuracy or completeness or use of the Omdia Materials. Informa TechTarget will not, under any circumstance whatsoever, be liable for any trading, investment, commercial, or other decisions based on or made in reliance of the Omdia Materials.

Get in touch: www.omdia.com askananalyst@omdia.com

