



Four crucial considerations when deploying edge servers



The edge has tougher requirements than traditional compute, but it also promises higher rewards to enterprises that approach it the right way.

For decades, the crown jewels of every enterprise's technology strategy lived in their on-premises data center. Then, as IT organizations adopted hybrid environments, big portions of their tech stack moved to the cloud. But today, the most impactful technology in many enterprises is neither on-premises nor in the cloud. It's at the edge.

According to IDC, nearly 80% of enterprises consider their edge workloads highly critical to the success of the business, and one out of every four servers they acquire by the year 2027 will be deployed at the edge.¹ But computing at the edge presents unique challenges not found in data centers or cloud operations. That's because the edge is increasingly where businesses' most important work takes place.

For example, edge devices inside factories use computer vision to identify defective products on assembly lines and ensure workplace safety. AI-enabled cameras in stores help retailers analyze customer behavior and reduce losses. Health monitoring systems in hospitals help doctors diagnose illnesses more quickly and heal patients faster.

¹ ["Gaining timely insights with AI inferencing at the edge,"](#) ITPro, Dec 2024.

In hundreds of other work scenarios across nearly every industry, edge devices are collecting data, analyzing it in real time, and driving insights. By 2029, 6 out of 10 edge deployments will be using composite AI solutions for generating content and making inferences.²

But because these devices must operate flawlessly across a broad range of environments, there is no one-size-fits-all approach to a successful edge deployment. Manufacturers will have different requirements for nontraditional environments and noise reduction than retailers, which in turn may have different concerns around physical and data security than healthcare providers. A big-box chain deploying thousands of edge servers will face significantly greater device management challenges than a local retailer deploying just a handful. And so on.

Because enterprises need flexible edge solutions that can adapt to a diverse range of environments and requirements, IT leaders need to consider the following four factors.

1. Environmental constraints

Edge computing devices are often deployed in environments that are dirty, dusty, and subject to abnormal temperatures — places like automobile assembly lines, oil rigs, cell towers, and fleets of semiautonomous vehicles, to name a few. These devices need to withstand temperature and vibration extremes not found in office, retail, or healthcare environments.

On the other hand, edge servers deployed in a public space such as a store, bank, or doctor's office don't need to be as physically robust. Instead, they must be smaller, consume less power, and produce less noise.

When designing edge architectures for the enterprise, IT leaders need to consider not only whether a server meets the organization's technology needs but also its optimal operating temperatures, air filtration requirements, decibel ratings, power demands, and other physical characteristics.



² ["Edge AI Success Demands Technology Curated for the Edge,"](#) VMware® Blogs, Oct 31, 2024.





2. Data and physical security needs

Organizations deploying servers at the edge need to worry about both the security of the data they're collecting and the devices themselves. The measures required to ensure physical and virtual security will often depend on where the machines are located and the types of data they are handling.

For example, servers accessible to industrial employees or those deployed in public settings like retail stores may be more vulnerable to physical tampering. They will require features like bezel locks and chassis intrusion kits to keep prying fingers out.

If edge devices are collecting sensitive data such as financial transactions, personally identifiable information, or medical records, they must be hardened internally as well. They will need integrated, enterprise-grade security measures such as self-encrypting storage, AI-powered intrusion detection, and a silicon root of trust that can prevent attacks that could compromise server firmware.

3. Performance requirements

Not all edge computing jobs are created equal. Server performance needs will vary greatly, depending on the use case. For example, retailers don't require a 64-core server with 256 gigabytes of RAM and a rack of GPUs to manage

store operations. But they may need that much computing horsepower to observe consumer behavior that lets them analyze foot traffic data to improve store layouts and ensure shelves remain fully stocked. They will likely want more powerful servers for their in-store video capture systems than those used for running their legacy virtual machines.

Nearly all organizations will need flexible systems that can scale as the size and sophistication of their edge deployments increase.

4. Management complexity

By 2030, an estimated 40 billion edge devices will be deployed globally,³ greatly increasing complexity for IT leaders across every industry. Manually patching or upgrading tens of thousands of edge devices in a timely fashion will soon become impossible for any single enterprise.

Remote management capabilities can reduce complexity and standardize practices across broadly distributed computing environments. Centralized consoles allow IT to manage all of an organization's edge devices from a single location and automate firmware updates and security fixes as its footprint expands. Remote management tools can also help tech shops detect and manage so-called shadow edge deployments, which emerge when individual business units stand up their own devices without input from IT.

³ "State of IoT 2024: Number of connected IoT devices growing 13% to 18.8 billion globally," IoT Analytics, Sep 4, 2024.





Engineered for work at the edge

No single server model can meet every requirement of every edge use case, but some come closer than others. Over the past year, Hewlett Packard Enterprise has worked with its largest customers to determine what they really need from their edge servers. The result is a server specifically engineered for the broadest possible range of deployments.

The HPE ProLiant DL145 Gen11 Server is the most versatile and flexible edge server HPE has ever created. With a footprint that is 50% smaller than the previous generation of servers, the HPE ProLiant DL145 Gen11 Server can sit on a tabletop or fit inside a wall or cabinet. It can withstand the dust, vibration, and extreme temperatures of an industrial environment, but it also runs cool and quiet enough to operate inside a store or a doctor's office. Enterprise-grade data and physical security features enable businesses to safely deploy the machines in high-traffic settings.

At the same time, a single HPE ProLiant DL145 Gen11 Server can handle up to 128 virtual machines, with support for the compute power needed to ingest huge volumes of data and deliver AI insights in real time. But perhaps its most compelling feature may be its simplicity: The HPE ProLiant DL145 Gen11 Server is built from the ground up to work with HPE Compute Ops Management, which allows for seamless remote operation at scale. Customers that have deployed HPE Compute Ops Management have slashed the amount of time devoted to managing servers by 75% while reducing downtime by an average of nearly five hours per server each year.⁴

As we move further into an AI-driven world, enterprise IT will need to adopt an edge-centric deployment strategy. The edge is where businesses will increasingly unlock new opportunities, and those that don't refine their edge strategy today are likely to struggle to catch up.

⁴ "New Technology: The Projected Total™ Economic Impact Of HPE Compute Ops Management," a commissioned study conducted by Forrester Research, June 2024. Results are for a composite organization.

Learn more at

HPE.com/ProLiant/DL145-gen11

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.

VMware is a registered trademark or trademark of VMware, Inc. and its subsidiaries in the United States and other jurisdictions. All third-party marks are property of their respective owners.

a50012028ENW