

HPE ProLiant Compute

eBook

Edge vs. cloud vs. hybrid: Strategic decision guide for CIOs and IT operations executives

Get started →

HPE





3 Executive summary

5 When to compute at the edge versus in the cloud

7 Industry perspectives

- 7 Healthcare
 - 8 Retail
 - 10 Manufacturing
 - 12 Telecommunications
 - 14 Distribution and Logistics
-

16 HPE ProLiant edge portfolio

19 Strategic guidance for executives


21 Take the next step

Executive summary

CIOs and IT operations leaders across industries face a common architectural dilemma: What belongs at the edge, what in the cloud, and how to integrate a hybrid strategy. Whether in healthcare, retail, manufacturing, telecommunications, or distribution, enterprises must balance latency, security, compliance, and cost when choosing where to run workloads. Despite differences in vocabulary and regulations, these industries share key trade-off criteria for edge vs. cloud vs. hybrid deployment models.

- **Latency and determinism:** Time-critical processes (real-time vision, robotics, telecom radio units, and so on) demand ultra-low latency and deterministic response, favoring on-site edge computing. Less time-sensitive tasks can reside in centralized cloud environments.
- **Data gravity and egress costs:** High-volume data sources (for example, video streams, sensor feeds) are expensive to backhaul. Preprocessing and analyzing data at the edge reduces cloud transmission needs, while cloud or core datacenters retain aggregated data for deeper analytics.
- **Bandwidth and intermittency:** Unreliable or limited WAN connectivity pushes critical functionality to the edge to avoid downtime. Edge sites must operate autonomously if network links go down.
- **Sovereignty and compliance:** Regulations around data locality (patient data in healthcare, consumer info in retail, and so on) often require data to stay on-premises or in-country. Edge and on-prem solutions ensure compliance with hybrid clouds for aggregation and auditing.
- **Environmental and physical constraints:** Remote and field locations face space, power, noise, and environmental extremes (dust, heat, vibration). These conditions necessitate ruggedized, compact servers at the edge, engineered for reliability outside pristine data centers.

[Read more >](#)



- **Security and resilience:** Distributed sites expand the attack surface so Silicon-rooted security and centralized policy control become crucial to protect edge nodes from tampering and cyber threats. Likewise, local processing ensures continuity if cloud access is cut off.
- **Operations at scale:** Managing hundreds or thousands of edge sites is impractical with one-off manual processes. Organizations need centralized, automated lifecycle management for provisioning, updates, and incident response across the edge fleet.
- **Economics and sustainability:** Power and cooling are often constrained at the edge, and network costs can skyrocket if all data is sent to cloud. An optimal mix (hybrid) **rightsizes** each workload's location to minimize total cost and carbon footprint.


Key takeaway: The winning architecture in all these sectors is hybrid—placing each workload where it best meets security, performance, and cost objectives—with edge computing handling real-time and local needs, and cloud/core handling centralized processing and aggregation. Modern edge platforms, such as HPE ProLiant servers, provide a consistent and secure compute foundation across this hybrid continuum, enabling enterprises to implement a unified strategy from edge to cloud.



When to compute at the edge versus in the cloud

Strategic drivers help determine which workloads belong at the edge (on-premises at clinics, stores, factories, cell sites, and so on), which are in the cloud, and which are in a hybrid pattern. The following framework (derived from the shared criteria above) guides CIOs in making those decisions:

- **Need for real-time processing:** If milliseconds matter—for example, a machine vision system on a manufacturing line rejecting defects, or a hospital ER dashboard analyzing vital signs—the compute should be on the edge for ultra-low latency and reliable response. Conversely, batch analytics or less urgent calculations can be in the cloud.
- **Volume of data vs. network costs:** Workloads generating massive data (HD video, IoT sensor floods) are ideal to handle at the edge to filter and summarize data locally. This avoids overwhelming WAN links and incurring high egress fees. Post-processed results or compact insights can then synchronize to cloud for aggregation.
- **Connectivity constraints:** Scenarios where sites might be disconnected or bandwidth-constrained (remote retail stores, distribution centers in rural areas, ships, or vehicles) require local edge computing for autonomy. Systems such as point-of-sale or inventory management must keep running even if the cloud is unreachable. In contrast, centrally managed corporate apps can live in cloud data centers when connectivity is assured.
- **Data residency and privacy:** Sensitive data subject to privacy laws (patient health records, retail customer PII) often must remain on-site or within the country. An edge or on-prem deployment meets compliance. The cloud can still be leveraged for anonymized data or global aggregations, but raw sensitive datasets stay at the edge for processing and short-term storage. [Read more >](#)



- **Physical environment:** If the workload must run in non-data center environments—for example, a factory floor, or a telecom base station hut—ensure the infrastructure is designed for those conditions. Here, specialized edge servers that are rugged (dust-filtration, extended temperature range, shock/vibration resistant) and compact are mandatory. Standard cloud servers would not survive or would be impractical in these locations. The cloud suits workloads in controlled environments, whereas the edge handles the harsh locales.
- **Security posture:** Edge sites outside the central data center’s physical security are more vulnerable to intrusion. If a workload runs at dozens of stores or clinics, each site should have platforms with built-in silicon security, encryption, and automated firmware verification to prevent breaches. The cloud provides strong security in centralized facilities, but when using edge, choose systems with equal security rigor (for example, silicon root of trust from HPE) to mitigate the wider attack surface.
- **Scale and IT staff:** Consider your operational capacity. Managing 100s of servers individually is infeasible. If a solution cannot be managed centrally or **lights-out**, it might gravitate to a cloud service instead. However, modern edge management tools now let a small IT team oversee fleets of distributed servers centrally (for example, HPE Compute Ops Management can patch and monitor thousands of edge nodes through one cloud-based console). This makes large-scale edge deployments viable where previously cloud might have been chosen for manageability.
- **Total cost and ROI:** Weigh the cost of cloud usage (ongoing OpEx, data transfer fees) versus deploying edge hardware (CapEx plus maintenance). Often, a hybrid model is most cost-efficient—for example, preprocess data at the edge to shrink cloud workloads and thus cloud costs. Also, edge computing can save networking costs and enable faster decisions, which has its own ROI (for example, preventing equipment failure through edge AI saves downtime costs).



Key takeaways: IT should adopt a hybrid mindset—use the above criteria to allocate each application to edge, cloud, or hybrid accordingly. In practice, this means pushing time-sensitive, data-intensive, or locality-required workloads to distributed edge servers, while leveraging cloud for aggregate analysis, long-term storage, and global coordination. Crucially, a consistent technology stack across edge and cloud (common management, security, and automation) is needed to tie it together.

Put each workload step where it delivers the best combination of security, performance, cost, and operational simplicity—with HPE ProLiant Compute providing a consistent platform across both edge and cloud. The rest of this guide explores how this applies to specific industries and what an optimal edge architecture looks like for each, along with the competitive landscape for edge solutions.

Industry perspective

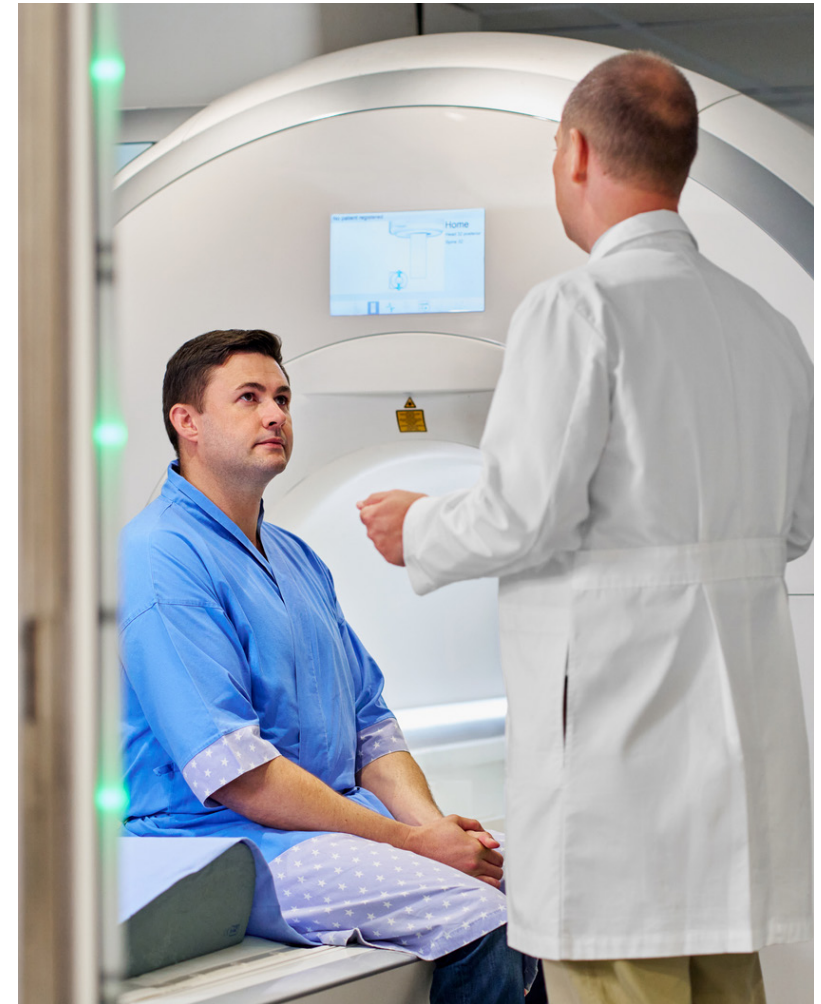
Healthcare: Smart hospitals and clinical systems at the edge

Hospitals and clinics generate vast amounts of data (imaging scans, patient monitoring, medical device feeds) that often require immediate analysis for patient care. Regulatory requirements (for example, HIPAA) also demand strict data privacy—often best met by keeping data on-premises. Edge computing in healthcare enables faster diagnosis and treatment, while ensuring sensitive data stays local.

Key edge use cases

- **Medical imaging diagnostics:** Radiology and pathology imaging can be analyzed on-site using AI models. For example, an edge server in a hospital can run AI inference on X-rays or MRIs as soon as they are taken, flagging anomalies within seconds.
- **Remote patient monitoring and telemetry:** In an ICU or a smart patient ward, edge systems can monitor vital signs streams and device telemetry (heart rate, blood pressure, ventilator metrics) continuously. By processing these signals locally, the system avoids any WAN latency and remains functional even if internet connectivity is lost.
- **Data locality for compliance:** Patient data often cannot leave the premises unencrypted. An edge infrastructure (on-premises data center or edge servers in hospital departments) ensures electronic health records, genomic data, and other personally identifiable information (PII) are processed and stored locally. For instance, a healthcare provider can use a hybrid cloud: An on-prem HPE server cluster runs the core hospital information system and analytics on patient data, while de-identified summary data is synchronized to a cloud repository for broader population health analysis.

[Read more >](#)



- **Clinical AI assistants:** Emerging AI applications such as automated ultrasound interpretation or surgery video analysis run at the edge in hospitals. This avoids relying on internet connectivity during critical procedures. For example, an AI video analysis of an endoscopy can run on an HPE edge server in the surgical suite, providing guidance to the surgeon in real time.

Edge advantage: Healthcare edge deployments reduce diagnosis times (which can save lives), preserve patient data privacy, and improve the reliability of clinical systems. By using rugged, compact edge servers (which can even be deployed in an emergency vehicle), hospital IT can bring compute power directly to medical staff.

Retail: Store-floor intelligence and resilience

In retail stores and distribution centers, connectivity can be spotty, and operations often cannot stop if the link to cloud is down. Retailers also gather in-store data—from video cameras, shelf sensors, POS systems—that is valuable for real-time operations and customer experience. Edge computing allows stores to remain operational independently and to leverage in-store data instantly for sales and loss prevention.

Key edge use cases:

- **Point-of-sale and inventory continuity:** By running POS systems, inventory databases, and pricing locally on an in-store edge server, stores ensure continuous operation. For example, a compact HPE ProLiant server at each store can host the store's sales system and sync to corporate cloud periodically. If the network goes down, the registers, card processing (often through local failover or edge payment appliances), and inventory lookups still work, avoiding lost sales.
- **Video analytics for shopper insights and loss prevention:** Retailers increasingly use camera feeds and AI at the edge. An edge server with GPUs can analyze live video to understand customer behavior (dwell times, heatmaps of foot traffic) and to detect theft or safety incidents in real time. These video analytics can also count customers for queue management or trigger restocking when shelves are empty—all processed on-site without needing to stream footage to the cloud. This reduces bandwidth usage and addresses privacy (video stays in-store).

[Read more >](#)



- **Personalized digital signage and experiences:** Edge computing enables dynamic in-store experiences. For example, a store could have an edge AI that recognizes when a loyalty customer enters (through their smartphone app or camera, if opted-in) and immediately tailors digital signage or offers for them. The edge server pulls that customer's preferences (stored locally or at edge cache) and updates screens in near-real time.
- **Omnichannel inventory and fulfillment:** For stores doing local fulfillment (buy-online-pickup-in-store), an edge server can quickly match online orders to in-store stock and direct staff for fast pickup prep. Speed and reliability here improve customer satisfaction.
- **Store environment management:** Edge IoT (though we avoid deep IoT discussion) still plays a role—sensors for refrigeration, HVAC, and so on, can be controlled by local edge logic to ensure compliance (for example, food safety) even if cloud commands are delayed. The edge system can raise alarms locally if a freezer temperature rises beyond threshold, triggering immediate fixes rather than waiting to notify a remote cloud.

Edge advantage: Retail edge deployments create resilient stores that keep transacting under any network conditions, and harness data-driven insights on the store floor. By processing data on-site, retailers reduce WAN costs (no need to stream every video to cloud) and can act on insights instantly.



Manufacturing: Smart factory and industrial edge

Factories, plants, and warehouses are highly instrumented environments where milliseconds of latency or a moment of downtime can mean scrap product or halted lines. Many manufacturing sites also generate enormous sensor and machine data volumes, and often operate in environments not friendly to traditional data center gear. Edge computing is central to Industry 4.0, powering real-time control systems, machine vision, and on-site AI for quality and predictive maintenance, all within the factory walls.

Key edge use cases

- **Computer vision for quality control:** Modern production lines use high-speed cameras and AI to inspect products for defects. These vision systems must process images instantly. Placing GPU-equipped edge servers on the factory floor (or in an adjacent control room) allows inferencing on each item as it comes off the line, without latency.
- **Predictive maintenance (IoT analytics):** Industrial sensors (vibration, temperature, pressure) generate high-frequency data. Instead of sending everything to the cloud, an on-site edge server can run ML models to predict failures in real time—like detecting abnormal pump vibrations in a steel mill and alerting maintenance before breakdowns. Local processing ensures fast response, even offline, and reduces data transfer costs through smart filtering.
- **Robotics and automation control:** Industrial automation relies on AGVs, robotic arms, and autonomous systems that demand real-time control. Cloud latency is too slow—local edge controllers (often rugged or integrated into HPE systems) coordinate robot actions in milliseconds. This ensures precise timing for motion control and safety, critical for synchronized operations on the factory floor.
- **Industrial AI and digital twin:** Manufacturers are using digital twins and AI at the edge to optimize machinery in real time. For example, an edge server can simulate a CNC machine and adjust settings on the fly for better performance and reduced wear. These models are compute-intensive, making powerful edge servers with modern CPUs or accelerators essential on the shop floor.



- **Harsh environment operation:** Many manufacturing sites have dust, extreme temperatures, vibration from heavy machinery, and so on. Regular rack servers might overheat or fail. HPE offers ruggedized HPE ProLiant servers built to MIL-STD specs for such environments. For instance, the HPE ProLiant DL145 Gen11 can operate from -5°C to 55°C and handle dust and vibration, making it suitable to mount near a production line. By deploying these, manufacturers can put compute closer to the action (for example, on a factory wall or control cabinet) rather than in a distant server room, reducing latency further.

Edge advantage: In manufacturing, edge computing boosts productivity, quality, and uptime by enabling real-time decisions that prevent costly defects and downtime. Many plants lack reliable, high-bandwidth connectivity, and even those that do can't depend on the internet for critical operations. Rugged edge systems designed for harsh environments provide self-sufficiency, with flexible deployment options like wall or cabinet mounting to suit factory conditions.




Telecommunications: 5G MEC and network edge

Telecom operators are building out multi-access edge computing (MEC) nodes and 5G edge sites to handle the explosion of data and low-latency service demands (think content caching, AR/VR, connected cars). Telcos require NEBS-compliant, compact servers at cell sites or central offices. They also often use open, cloud-like network functions (vRAN, mobile core) that run on standard servers, but at the network edge. Edge computing in telco reduces backhaul load and enables new services by processing data close to subscribers.

Key edge use cases

- **vRAN/O-RAN (virtual radio access network):** Traditionally, baseband processing for cell towers was done with proprietary hardware. Now, telcos are virtualizing RAN on standard x86 servers located at cell sites or aggregation points. These servers must handle extremely time-sensitive signal processing. HPE ProLiant DL110 Gen11, for example, is a telco-optimized 1U server designed for open RAN workloads; it can live at a 5G site and process radio signals with low latency. By deploying edge servers for RAN, operators gain flexibility and can allocate resources dynamically (and even share infrastructure among multiple purposes).
- **MEC for content and applications:** Telcos deploy edge computing in metro data centers or even on-premises at enterprises to offer services such as content delivery networks (caching popular video content closer to users), multiplayer gaming platforms with low lag, and industrial IoT application hosting.
- **Network analytics and optimization:** Telecom networks generate immense telemetry: traffic patterns, signal quality metrics, and user mobility events. Rather than shipping all logs to a central cloud, local edge analytics can occur. An AI model at the edge might analyze cell congestion in real time and dynamically adjust network parameters or offload users to less-loaded cells. This distributed intelligence improves network performance (users get better speeds and fewer drops) and reduces data that must be sent to central nodes.

[Read more >](#)



- **Private 5G and edge for enterprises:** Telcos (and partners) also deploy on-premises edge solutions for enterprise sites (such as a smart port or a large factory) as part of a private 5G network offering. These edge installations often run on the same hardware that telcos use—rugged servers that can handle both the network functions and the enterprise’s edge applications.
- **Telco cloud and NFV at edge:** Many network functions (firewalls, routers, CDN nodes) are virtualized and placed at strategic edge locations. By processing video closer to users, telecom operators can lower latency and reduce backhaul traffic. For example, Ericsson and GSMA report that edge-based video optimization and caching can significantly reduce bandwidth on backhaul links, improving user experience and lowering costs by adapting streams locally.¹

Edge advantage: Edge computing is essential to 5G, enabling low latency, high bandwidth, and new revenue opportunities like hosting customer workloads and powering IoT. It also reduces network costs by localizing traffic. Telco deployments require carrier-grade hardware with features like NEBS compliance, DC power, and high availability. Rugged edge systems built for telco environments offer short depth, filtered cooling, and precise timing support—ensuring secure, scalable management across distributed sites.



¹ Ericsson Edge Computing Whitepaper; GSMA Telco Edge Cloud Report




Distribution and logistics: Real-time supply chain at the edge

Warehouses, logistics hubs, and transportation fleets operate in real time and often in locations with limited connectivity (such as ships or rural logistics centers). Edge computing helps coordinate complex physical operations (inventory handling, fleet routing) with minimal latency and continued operation even when cloud connectivity is unavailable. It also enables safer, more efficient operations by processing data from vehicles, drones, or conveyors on-site.

Key edge use cases


- **Warehouse automation:** Modern distribution centers rely on robots for picking and inventory movement, requiring on-site edge control for real-time coordination. Local systems manage inventory, guide robots, and ensure uninterrupted operations—even during connectivity issues. Edge computing also powers visual AI tasks like verifying packed orders instantly, integrating barcode scanners, sensors, and robotics with sub-second responsiveness.
- **Fleet management and routing optimization:** Logistics companies use edge computing at depots to analyze vehicle telemetry, plan maintenance, and optimize routes in real time—without relying on cloud connectivity. Local systems can also process traffic and weather data to update drivers instantly. In mobile scenarios like shipping or air cargo, edge systems onboard vehicles analyze journey data continuously and sync when connectivity is available.
- **Real-time tracking and asset management:** For distribution, knowing the real-time location and status of inventory and assets is gold. Edge solutions in logistics hubs track every pallet via RFID, every forklift's position, and so on, and can generate a live **digital twin** of operations. The edge system can, for instance, immediately spot if a pallet is stored in the wrong location and alert staff before it causes a shipping error. These computations happen locally to handle the sheer volume of sensor pings in a warehouse and to respond without the latency of cloud.

[Read more >](#)



- **Safety and monitoring:** Similar to public safety, but in an industrial setting—cameras and sensors in facilities can detect safety incidents (like a worker entering a restricted zone or a conveyor jam) and halt machinery instantly. Here, edge computing with AI vision ensures an immediate response.
- **Environmental tolerance:** Distribution centers can be not climate-controlled, and full of dust from packaging. Thus, like manufacturing, the servers running on-site need to be robust. HPE offers rugged edge servers that can even be placed on a simple rack in a warehouse floor (some are half the size of a traditional server and can operate in high-temperature, dusty environments). This avoids needing a dedicated server room with AC in each facility, saving cost and complexity.







Edge advantage: The supply chain moves faster and with more intelligence through edge computing. Various companies pioneered heavy use of edge in warehouses, but now even mid-sized firms can deploy affordable edge servers to get similar benefits. The operational impact is huge: better throughput, fewer errors, and real-time visibility. And critically, distribution operations do not grind to a halt due to a network issue; local systems keep conveyors and scanners running.



HPE ProLiant edge portfolio

To power these industry-specific 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. Key HPE differentiators include ruggedized designs for harsh sites, embedded silicon-based security, and unified cloud-native management—all of which address CIO concerns about deploying and operating distributed edge infrastructure.

Table 1. HPE ProLiant Compute—edge server portfolio

Servers	Server model	Form factor	Optimized for	Industries
	MicroServer Gen11	Ultra-compact	Local IT services, including web hosting, email, and point-of-sale systems	Small business, hospitality
	ML30 Gen11	Mini tower	Enterprise-grade AI capabilities for SMBs and distributed sites	SMB, healthcare clinics, remote offices
	DL20 Gen11	Compact 1U	Foundational IT services such as file, web, and email hosting; POS analytics	Retail, SMB, branch offices
	DL110 Gen11	1U ruggedized	Virtualized and Open RAN architectures; network edge workloads	Telecommunications, service providers
	DL145 Gen11	2U ruggedized	AI inferencing, smart surveillance, industrial automation	Retail, energy, public sector, and manufacturing
	EL8000 / EL8000t	Modular, systems	Mission-critical applications; pull-replace-repair workflows; high availability	Public sector, industrial IoT Telecommunications

[Read more >](#)



This breadth means CIOs can choose the right node for each site: For example, a retail chain might use HPE ProLiant MicroServer in very small shops, HPE ProLiant DL20 1U racks in supermarkets, and a few HPE ProLiant DL145 servers at regional hubs for AI processing. Importantly, all these servers share common management and security—they run HPE iLO firmware with silicon root of trust, and they can all be administered through HPE Compute Ops Management platform.

Consistent management and security: A huge concern in edge deployments is how to manage and secure thousands of distributed servers with limited IT staff on-site. HPE addresses this in two ways:

- **HPE Compute Ops Management:** This is a cloud-native management platform that admins can use to oversee the entire server fleet (edge and data center) from a single console. Using HPE Compute Ops Management, IT can deploy firmware updates, monitor health, and even perform fast provisioning of new edge servers. The result: Up to 75% less time managing remote servers and significant savings in travel and manual effort, as found in a Forrester study.² For a CIO, this means you can scale your edge without proportionally scaling headcount. A small team can manage, distributed servers, pushing policies and configurations through unified console.
- **Silicon root of trust and built-in security:** Every HPE ProLiant edge server embeds silicon root of trust in its HPE iLO management chip. This means from the moment the server powers on, the firmware verifies itself against an immutable silicon fingerprint. It prevents compromised firmware from running—a critical defense since edge sites could be physically accessed by malicious actors. HPE is the only vendor that designs its own baseboard management controller (BMC) silicon for this purpose; competitors typically use off-the-shelf chips. Additionally, HPE servers offer physical security options such as chassis intrusion detection, locking bezels, and Kensington locks to deter tampering. For CIOs, this layered security means you can trust deploying servers in less secure locations. Even if an edge server is stolen or modified, the chain-of-trust (with secure boot, encrypted drives, and so on) protects your data and credentials.



[Read more >](#)

² Commissioned study conducted by Forrester Research, New Technology: The Projected Total Economic Impact of HPE Compute Ops Management, June 2024.




Performance with GPUs for edge AI: As highlighted in the industry cases, AI at the edge is a big trend (from vision to machine learning). HPE ProLiant servers are purpose-built to accommodate accelerators. For instance, the DL145 Gen11 supports up to three single-width GPUs (such as NVIDIA® L4) in a short-depth chassis, something quite unique for an edge form factor. The portfolio includes both Intel® and AMD-based systems to suit different workloads.

Rugged, remote environments design: Several models (DL145 and EL8000 notably) are engineered for diverse environments. For example, the EL8000 is a modular system that can be deployed in telecom cabinets, featuring high shock/vibration resistance and front-access modularity for quick swapping of server blades. It has been used in defense and telecom where dust filtering and -48V DC power are needed. The DL145 Gen11, as another example, is about 50% shorter in depth than a DL365 rack server and can be mounted on walls or under benches—it is quiet enough (~55 dB) for office environments too. These considerations matter; they allow IT to deploy compute in places previously impossible (like a manufacturing line's control enclosure or a retail backroom without dedicated cooling).

From a strategic standpoint, HPE's approach is to give a unified experience **from factory floors to hospital wings**, delivering the same manageability and security you expect in a data center, but in edge-optimized packages. This aligns well with the needs of IT ops execs who want to extend corporate IT to the edge without creating a separate silo of hard-to-handle devices.

For CIOs and IT executives, the choice may come down to priorities: If security and unified management are paramount—especially with lean IT staff—the HPE ProLiant edge computing portfolio provides a compelling solution with its silicon root of trust and cloud-managed simplicity.



Strategic guidance for executives

Here are strategic recommendations for CIOs and IT ops leaders planning edge deployments in the mentioned industries:

- **Embrace hybrid, but design for autonomy:** Plan every edge deployment to function on its own when needed. In healthcare, that means local data processing for critical apps; in retail, offline-capable store systems; in manufacturing, fail-safe controls on-site. Use cloud to augment, not to depend. The technology exists (as we have seen with edge servers and local AI) to make sites self-reliant while still reaping cloud benefits through periodic sync. Design your architecture such that edge and cloud complement each other—for example, edge handles immediate decisions, cloud handles global training or coordination.
- **Leverage edge AI for quick wins:** Identify one or two high-impact AI use cases in your operations and pilot them with edge computing. For example, a retailer might start with an AI loss prevention camera system in a few stores, running on an edge server; a manufacturer might implement AI-based quality control on one line. Use these to build expertise and ROI proof points. By keeping the AI at the edge, you ensure data privacy and low latency. As stats show, organizations see improved efficiency and resilience as top benefits of edge investments.
- **Choose platforms with unified management:** The operational burden can explode if each site requires hands-on IT. Select an edge platform that integrates with your core management tools or, better, offers a single pane for all edges. HPE Compute Ops Management is one example that can manage thousands of servers; The key is to budget and plan for management software and central teams (or outsourced NOC) to oversee the fleet. Factor these into TCO—sometimes paying a bit more for a platform with automated management can reduce costs.

[Read more >](#)



- **Prioritize security and zero trust at edge:** Treat every edge node as untrusted until proven otherwise. That means enabling secure boot, encryption for data at rest (self-encrypting drives; HPE offers those with root of trust tie-in), and network security (VPN/SD-WAN) for communications. Consider vendors with proven supply chain security (HPE servers, for instance, are built with a supply chain attestation—tamper-free during transit—which might align with your compliance needs). Implement central logging and monitoring of edge devices—many breaches in recent years exploited neglected edge systems. The good news: Hardware such as HPE can enforce a lot of this out of the box, and tools exist to maintain uniform security policies across sites.
- **Plan for scale and future-proofing:** Edge deployments tend to start small (a few sites) but can rapidly expand if successful. Ensure your chosen solution can scale to many sites, and will support future needs like more AI, more storage, or integration with new technologies (such as 5G slicing for enterprise, or upcoming regulatory changes requiring on-site data retention). A modular system (such as EL8000 from HPE for incremental growth, or ability to upgrade CPUs/GPUs in place) might pay off. Also, consider power and space: the industry is predicting that AI will become a dominant driver of edge infrastructure investments in the next few years, so you might need to add GPUs or accelerators to existing edge servers. Choose platforms that have headroom (PCIe slots, power budget).
- **Evaluate as a service models:** Many leading vendors now offer edge infrastructure in an as-a-service model. HPE GreenLake cloud can deploy and manage edge servers with consumption-based pricing, which might appeal if you want to avoid owning equipment or if workloads are variable. Dell has Apex and Lenovo TruScale similarly. This can reduce up-front cost and allow the vendor to handle some management—but ensure service-level agreement (SLA) specifics, especially for remote sites (for example, who replaces a failed edge server under HPE GreenLake, and how fast?). For distributed retail or branch scenarios, as-a-service can significantly simplify procurement and lifecycle management as you refresh hundreds of small sites.

- **Pilot in each industry context:** The needs of a hospital vs. a factory vs. a retail store differ; it is wise to pilot edge solutions in each context rather than a one-size-fits-all approach. You might find, for instance, that a rugged 2U server is overkill for a retail store but perfect for a factory. Use pilots to adjust configurations—maybe try built-in filtration design is needed in a dusty warehouse site, or maybe in telecom, you need DC power options. HPE can often assist with proof of concepts; take advantage to test things like latency improvements and failure scenarios.

By following these guidelines, CIOs can confidently extend their IT footprint to the edge—delivering new capabilities on-site while maintaining central oversight. The technology and solutions discussed (particularly HPE ProLiant portfolio with its edge-centric innovations) show that it is entirely feasible to run enterprise-grade workloads outside the data center. In doing so, organizations in healthcare, retail, manufacturing, telco, and distribution can gain agility, improve customer experience, and drive operational efficiencies that directly impact the bottom line.

Final thoughts

The edge is no longer a fringe concept—it is a strategic component of modern IT architectures. The decision is not edge vs. cloud as an either/or, but rather aligning each workload to its optimal placement and having a hybrid strategy to tie it together. With the right edge infrastructure, executives can deliver real-time intelligence and services at the point of need (be it a bedside, checkout lane, assembly robot, cell tower, or loading dock), all while keeping control, security, and insight through their centralized tools. This guide has outlined how to approach that journey, and how HPE ProLiant solutions in particular can support it, along with an understanding of alternatives. Armed with this, CIOs and IT ops leaders can make informed decisions to thrive at the edge.





Take the next step

Whether you are modernizing clinical diagnostics, optimizing retail operations, automating factory floors, scaling telecom infrastructure, or streamlining logistics, the right edge strategy can achieve real-time intelligence, resilience, and competitive advantage.

- Ready to evaluate your edge architecture? → Start with a pilot in one high-impact site and validate performance, security, and manageability.
- Need help choosing the right edge platform? → [Explore HPE ProLiant edge portfolio](#) and request a tailored recommendation based on your industry and workload needs.
- Want to simplify operations across hundreds of sites? → [Learn how HPE Compute Ops Management can centralize lifecycle management](#) and reduce operational overhead.
- Considering a consumption-based model? → Discover HPE GreenLake for edge and see how you can scale edge infrastructure with cloud-like flexibility.

Learn more at

[HPE.com/proliant/edge-computing](https://www.hpe.com/proliant/edge-computing)

[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 is a trademark of Advanced Micro Devices, Inc. Intel is a trademark of Intel Corporation or its subsidiaries in the U.S. and/or other countries. NVIDIA is a trademark and/or registered trademark of NVIDIA Corporation in the U.S. and other countries. All third-party marks are property of their respective owners.

a00151462ENW

HEWLETT PACKARD ENTERPRISE

[hpe.com](https://www.hpe.com)

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

