WHITE PAPER

CHATSWORTH PRODUCTS

The Evolving Edge: Designing the Edge Infrastructure for the Next Wave of Distributed Intelligence

By
Michael Moore
Sr. Product Manager, Open Architecture Systems

David Knapp
Product Marketing Manager

Published: February 2026

chatsworth.com

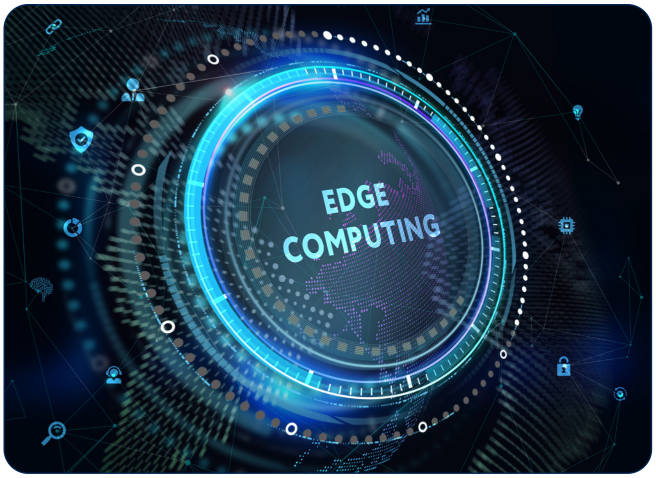

Edge computing has become the operational backbone of modern IT. Across industries—from manufacturing and healthcare to utilities and retail—what began as pilot projects to reduce latency or offload cloud traffic has matured into a massive, distributed ecosystem powering real-time decisions, automation, and AI workloads at the point of action.

Traditionally, "the edge" meant a physical location — a remote site, a branch office, a warehouse, a factory floor, a retail store. It was shorthand for: *anything not in the core data center.*

But today's edge is more than a location—it's a distributed intelligence network.

At the modern edge, instead of sending everything to a central data center, compute, analytics, and decision-making is moving closer to where data is created, allowing each location to act right away, not just storing and forwarding data, but thinking with it. At CPI, we define this "new edge" as the intersection of **AI-driven performance, physical resilience, and environmental responsibility**—where reliability, not proximity, determines success.

For industries such as retail, healthcare, manufacturing, and transportation, this edge-centric approach is as much a financial strategy as a technical one. By processing data closer to where it's generated, organizations can control cloud egress costs, improve AI performance, and achieve greater predictability in both cost and compute efficiency.

According to IDC, global spending on edge computing solutions accounts for nearly $261 Billion in 2025 and is projected to grow at a compound annual growth rate (CAGR) of 13.8%, reaching $380 Billion by 2028, underscoring how the edge has shifted from *proof of concept to critical dependency.*[1]

The convergence of **AI, IoT, and operational technology (OT)** has transformed the edge from a simple data-collection layer into thousands of miniature data centers—each with its own thermal, electrical, and environmental complexity. In traditional data centers, reliability is ensured through climate control, redundant power, and on-site staff. At the edge, operators must deliver those same outcomes—without those advantages.

Global Edge Spending by Technology Groups

(2025 forecast)

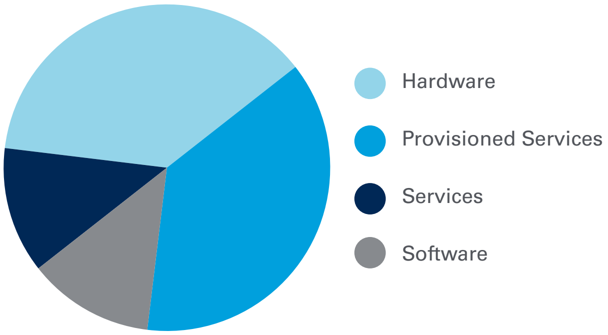

Source: IDC Semiannual Edge Spending Guide, Forecast V1 2025

When every location becomes a data environment, the physical layer is no longer background infrastructure. It becomes the foundation of performance, efficiency, and uptime.

At Chatsworth Products (CPI), we view this shift through a focused lens: **engineering the physical infrastructure that turns edge ambition into operational reliability.**

This white paper explores what success looks like at the new edge—one defined not by location, but by resilience engineered into every enclosure, cabinet, and connection.

The Transformation Drivers: How Data, AI, and Convergence Are Reshaping Edge Operations

To understand how physical edge infrastructure needs to evolve, we first need to examine what's driving the transformation—the technologies, market forces, and operational realities that are redefining the role of the edge.

Rapid Scale: From Pilot Projects to Production Networks

- **Edge deployments are expanding rapidly—**from a few test sites to hundreds of production locations. IDC projects $381 billion in spending by 2028, a 13.8% CAGR from 2025.[2], and those deployments increasingly support mission-critical functions—from hospital imaging to autonomous logistics.
- **Because use cases are diversifying, infrastructure must reliably support highly varied use cases—**retail (inventory, personalization), autonomous vehicles and mobility, healthcare (diagnostics at point of care), industrial automation—each with distinct environmental, latency, and compliance demands. This complexity amplifies the risk profile of every design choice.
- **The new reality:** The old edge could tolerate variability - the new edge cannot. Achieving scale at the edge isn't about adding more sites—it's about deploying them identically, securely, and repeatably. Small inconsistencies in cable management, airflow, or power distribution that were negligible in a few racks now multiply into systemic reliability risks when replicated across global deployments.

Old Edge vs. New Edge

Characteristic	Early Edge	The New Edge
Primary Role	Data collection and network aggregation	Real-time AI processing and automation
Compute Density	Low (<1 kW/cabinet)	High (5–10 kW + per cabinet)
Environment	Controlled or conditioned	Uncontrolled, distributed, and remote
Reliability Model	Periodic maintenance	Continuous, autonomous monitoring
CPI Infrastructure	Standard network racks	Intelligent, thermally managed enclosures (e.g., VersaEdge™ Wall-Mount Cabinet, RMR®)

The Convergence of AI, IoT, and 5G: From Connected Devices to High-Density Compute

The first generation of edge systems connected sensors and sent data to the cloud. The new generation processes that data locally, running **AI inference, analytics, and server-intensive workloads** directly at the edge.

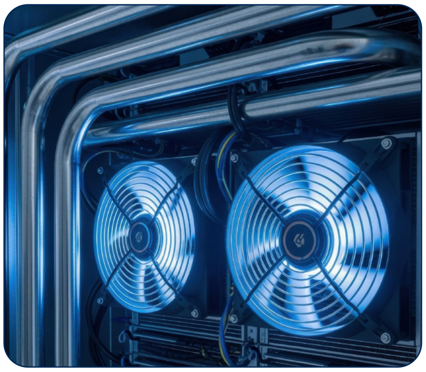

This convergence of **AI, IoT, and 5G** has transformed the physical requirements of edge sites. AI workloads create **heat densities exceeding 6 kW,** unpredictable power spikes, and constant data throughput. Meanwhile, 5G's low-latency expectations mean any performance dip, whether caused by heat or power fluctuation, directly affects service quality.[3]

The new reality: The edge is no longer low power, rather **high-density compute in miniature**. Reliability now depends less on enclosure ruggedization and more on predictable thermal performance. Modern enclosures must integrate **defined airflow paths, active fan systems, and sealed cable ingress** that maintain temperature control in unconditioned spaces.

At CPI, decades of passive cooling innovation—refined in hyperscale environments—now inform the design of our edge-ready cabinet ecosystems. By integrating containment-grade airflow and optional hybrid liquid cooling within compact enclosures, CPI extends **data-center-level thermal predictability** to the distributed edge.

Industry	Real-World Applications [4]
Retail	Frictionless checkout, inventory tracking, and store analytics
Media & Entertainment	Video tagging, encoding, and regulation compliance
Telco	Modernization of networks for 5G Core, RAN, and local data processing
Financial	Data sovereignty, low-latency trading, and payment processing
Healthcare	AI for imaging diagnostics and patient monitoring
Manufacturing	Quality control, IP and regulations, and autonomous mobile robotics
Agriculture	AI-equipped drones for real-time insights in crop management
Education	AI integrated into provision headsets for on-the-spot inferencing
Transportation	Data collection, autonomous driving systems safety

Full Forms & Terminologies to Remember [5]

AI (Artificial Intelligence)
Machines that can learn, reason, and make decisions.

IoT (Internet of Things)
A network of connected devices that collect and share data.

AIoT (Artificial Intelligence of Things)
The integration of AI with IoT to create intelligent, automated systems.

Regulatory and Sustainability Pressure: From Connectivity to Accountability

The first wave of edge expansion focused on connectivity. The next is being shaped by **efficiency, transparency, and accountability.**

As governments and enterprises tighten environmental and cybersecurity standards, operators must now prove that every site meets energy efficiency, uptime, and security requirements. Without integrated monitoring and data-driven reporting, those validations are impossible and inefficiencies hidden across hundreds of sites can compound into significant cost and compliance risks.

As edge infrastructures proliferate, small leakages in reliability compound across many sites, impact both for regulatory compliance and operating expenses.

The new reality: Sustainability and performance are now inseparable. Intelligent infrastructure, equipped with **remote sensors, power monitoring, and access control,** is essential not only for operational efficiency but also for meeting regulatory and Corporate Social Responsibility (CSR) mandates.

CPI's approach embeds sustainability at the material and operational levels—using **U.S.-based manufacturing,** recyclable materials, and **integrated monitoring** to help customers document energy efficiency, uptime, and CSR compliance across all sites.

Resource Constraints:
From Staffed Facilities to Autonomous Operations

Edge sites are smaller, more numerous, and increasingly unstaffed. Power quality may fluctuate, ambient conditions may shift, local expertise limited, and downtime response is measured in hours of travel rather than minutes of intervention. Operators are being asked to do more—with fewer hands and less margin for error.

Every unplanned site visit or downtime event eats into both OPEX and customer confidence. Small inefficiencies are amplified at the edge, turning infrastructure design into a critical business risk, not merely an engineering challenge.

The new reality: Today's edge enclosures must function more like an autonomous system than a remote outpost. Infrastructure must not only protect equipment but also provide environmental telemetry, automated alarms, and secure access control to maintain uptime without human oversight.

Together, these forces define the new edge. **Physical infrastructure is no longer background support; it is the operational backbone.** Its evolution will determine how far edge technology can advance—and how effectively operators can scale intelligence without multiplying risk.

Redefining Edge Success:
Why Physical Edge Infrastructure Must Evolve Now

From these pressures—rapid scale, AI convergence, sustainability mandates, and limited resources—a clear truth emerges: **The edge is no longer defined by where data is processed—but by how well it performs, anywhere it's deployed.**

Software defines what the edge can do. Physical infrastructure determines whether it can sustain that capability—how effectively it is powered, cooled, secured, and maintained under real-world conditions.

Downtime tolerance has all but disappeared. Each localized failure now carries network-wide consequences:

- In **manufacturing,** an hour of lost machine data can derail production.
- In **healthcare,** edge systems process real-time diagnostics and imaging.
- In **utilities and telecom,** local compute underpins critical service continuity.

In this environment, physical infrastructure is no longer passive support—it's a determining factor in operational continuity. Managing hundreds or thousands of distributed nodes requires systems engineered for **data-center-grade predictability** within a fraction of the footprint and often in far harsher conditions.

Achieving that level of predictability requires a new approach to physical design—one built around three core infrastructure priorities that define performance at the distributed edge.

Designing for Distributed Intelligence: Three Infrastructure Priorities for Predictable Performance

Bringing data-center reliability to the edge requires more than ruggedized hardware — it demands **precision engineering across the physical layer.**

Together, these elements determine whether an edge network performance under harsh environmental conditions will have a competitive advantage, especially in industries where reliability is non-negotiable.

1. Thermal Control: Not Just for Data Centers Anymore

Explosive growth in 5G, IoT, and AI workloads pushing thermal limits at the edge. Power-dense servers and accelerators can quickly exceed safe operating temperatures, leading to throttling, unreliable inference performance, and premature hardware wear. Thermal issues can't be something engineers deal with later. They need to be addressed at the design stage.

"The diversity of edge environments means that a one-size-fits-all cooling approach doesn't work. In fact, some of the best thermal strategies start with environmental mapping. Before designing a cooling solution, developers need to understand the conditions under which their devices will operate." [6]

"Additionally, the cost-sensitive nature of edge operations increases reliance on efficient airflow management to maximize ROI. Containment solutions help edge facilities achieve enterprise-grade reliability despite their smaller size. With the global rollout of smart cities and 5G, demand for edge containment will only accelerate. This positions edge data centers as the fastest-growing type in the containment market." [7]

We're now seeing a shift toward more aggressive and proactive thermal approaches:

CPI leverages its decades of expertise in passive and hybrid cooling to deliver data-center-grade containment principles to VersaEdge™ Wall-Mount Cabinets, enabling up to 10 kW of passive cooling.

- **Optimized Airflow Designs:** Enclosures designed to promote smooth air circulation, preventing hot air from getting trapped around equipment. This helps maintain stable inlet temperatures and avoids "hot spots" that can lead to overheating.
- **Integrated Fan Kits** pull hot air out and push cool air through the enclosure. This turns the wall-mounted cabinet into a mini thermal ecosystem, capable of supporting denser compute equipment like edge servers or switches.
- **Modular Depth Options to accommodate different thermal profiles**. Larger depths provide more room for airflow and accessory cooling kits, which is essential for edge compute with higher wattage.

VersaEdge™ Wall-Mount Cabinet air flow

By applying airflow containment science, CPI helps operators support up to **10 kW per cabinet passive application with no supplemental fans needed.**

2. Power Visibility and Control from a Distance

At the edge, power failures aren't just outages — they're extended blind spots. With limited or no on-site staff, operators need power systems that can be diagnosed, controlled, and corrected remotely. According to Uptime Institute, **power failures cause nearly 54% of all IT service outages.**[8]

"In distributed edge deployments with limited or no on-site staff, even a single unmonitored or malfunctioning circuit can lead to prolonged downtime lasting hours or days due to the difficulty and delay in physical intervention" - Uptime Intelligence Annual Outage Analysis Report 2025 [9]

Annual Outage Analysis

(2024 survey)

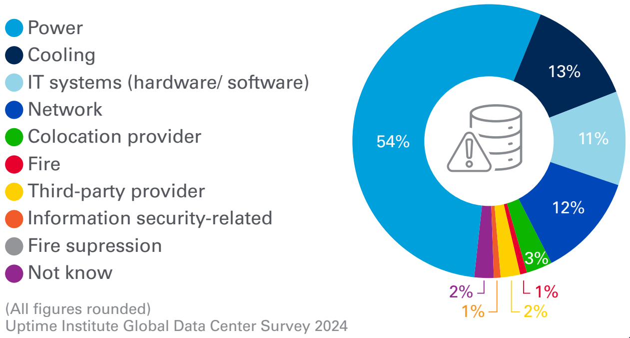

(All figures rounded)
Uptime Institute Global Data Center Survey 2024

Uptime at the edge isn't just about redundancy—it's about **situational visibility.** Operators need real-time insight into voltage quality, breaker coordination, and load balance to act before a failure occurs. Power management must evolve from reactive response to **proactive control**.

Predictive power management—driven by intelligent metering and data correlation—is redefining uptime strategy. The industry is shifting from reactive response toward proactive, predictive power management with power systems that combine redundancy, intelligent metering, and remote control to maintain reliability without physical intervention.

Power Capabilities Required for Edge Autonomy

CPI eConnect® PDU with Quadlock outlets

- **High-accuracy metering** at the inlet, branch, and outlet level identifies anomalies before they become outages.
- **Remote switching and access control** enable technicians to restore functionality or isolate issues without dispatching a field team.
- **Optional environmental sensor ports** allow power and environmental data to be correlated—turning energy data into actionable intelligence.

The result: faster root-cause analysis, fewer truck rolls, and the ability to maintain uptime across hundreds of sites with a lean operational footprint.

Intelligent Power Distribution Units (PDUs) form the foundation of this shift. CPI's **eConnect® PDUs** provide enterprise-grade power visibility and remote control across hundreds or thousands of edge sites from a centralized interface. Granular metering, remote outlet control, and integrated sensor support enable operators to move from reactive troubleshooting to predictive intervention, identifying risk conditions early and resolving issues without site visits.

3. Condition-Aware Infrastructure: Environmental and Access Control

Edge environments can no longer be treated as passive containers for IT equipment. As infrastructure pushes into remote, unmanned, and thermally unstable locations, the physical infrastructure layer within the enclosure becomes an active part of system reliability. If conditions inside and around the cabinet aren't visible, the edge site is effectively operating without situational awareness.

"Studies indicate that up to 60% of unplanned outages at distributed edge sites originate from physical and environmental factors such as heat, humidity, dust ingress, and unauthorized access, rather than compute hardware faults" [10]

At distributed edge sites, failures rarely begin with servers or switches. They begin with rising inlet temperatures, moisture exposure, particulate contamination, or an open door that goes unnoticed. These risks develop quietly at the physical layer, outside the reach of traditional IT monitoring, until performance degrades or systems shut down.

That's why environmental sensing and access monitoring are not optional enhancements—they are foundational requirements for predictable uptime. Application alarms confirm failure after it occurs.

True edge autonomy is achieved when each enclosure is equipped with sensors and control mechanisms that continuously monitor environmental conditions and access activity, reporting anomalies in real time—while corrective action is still possible.

Condition Awareness at the Physical Layer Requires:

- **Continuous environmental sensing** within the cabinet environment, including temperature and humidity at key intake and exhaust points
- **Verified physical access awareness,** with logged and time-stamped door events
- **Unified visibility,** where power, environmental, and access data are correlated in a single operational view

Condition awareness only delivers value when these signals are collected at the enclosure itself and surfaced in a way operators can act on. Siloed sensors or disconnected logs create data, not insight.

That's why intelligence must start at the physical layer closest to the equipment. Sensors and access controls operate within the enclosure environment, capturing local conditions and events and feeding them into centralized monitoring platforms where anomalies can be identified early and addressed remotely.

CPI supports this approach in the following ways:

- **Integrated sensor ports** on eConnect® PDUs support plug-and-play temperature and humidity probes at critical intake and exhaust points and soon you will be able to monitor seamlessly with the eConnect® Sensor Array.
- **Electronic locking systems** record and timestamp access attempts, enabling audit trails and alerts for physical security events.

When paired with **Power IQ® for eConnect,** operators gain a **single-pane view** of power, environmental, and access conditions across every distributed site.

The result: every edge location becomes **an observable system**, not a black box. Operators can monitor, correlate, and respond to conditions in real time—transforming environmental data into actionable insight.

Integrated sensor ports on eConnect® PDU

Electronic locking system

Power IQ® for eConnect Dashboard

Scaling the Edge: A Playbook for Repeatable, Modular Physical Infrastructure

Getting thermal performance, power integrity, and environmental visibility right isn't the finish line. Scaling the edge is not about designing the perfect enclosure once—it's about designing a **system that performs consistently everywhere.**

Key Ingredients for Scalable Physical Edge Enclosures:

CPI's approach distills physical design into three disciplines— edge and remote readiness, scalability through modularity, and governance at scale —each ensuring that distributed intelligence can scale with speed, predictability, and control.

1. Edge and Remote Readiness

At scale, environmental variability is the first source of operational inconsistency. When deployments span unconditioned rooms, outdoor shelters, and production floors, infrastructure must be designed to perform predictably **without site-specific engineering or manual oversight.** Operational assurance comes from enclosures that establish a stable, self-contained operating environment regardless of external conditions. NEMA- and IP-rated architectures, sealed cable ingress, filtration, and optional hybrid cooling allow operators to deploy into thermally volatile or unmanned locations without redesigning the solution each time.

2. Scalability Through Modularity

Once performance is stabilized across environments, the next barrier to scale is physical variation between sites. Edge sites vary in wall depth, airflow availability, mounting conditions, and cable routing—but treating each variation as a new engineering effort quickly makes scale impossible. Modularity **built around a stable core design** delivers consistent structural, thermal, and power behavior, regardless of where it is deployed. Standardized components and repeatable design patterns reduce lead times, minimize installation risk, and create predictable cost envelopes across deployments. Logistics are engineered into the system itself, with form factors that fit standard doorways, adapt to uneven mounting surfaces, and arrive preassembled for rapid, low-touch installation.

3. Governance at Scale

Scalable edge platforms rely on **pre-validated security and compliance frameworks.** Enclosures must ship with consistent locking mechanisms, rating certifications, and audit-ready access controls that reduce the need for site-by-site customization or post-install remediation. This allows organizations to deploy edge infrastructure rapidly while maintaining uniform governance standards.

By integrating security and compliance directly into the physical design—through rated enclosures, standardized access control options, and documented certification—CPI enables operators to scale distributed deployments without introducing operational exceptions or compliance gaps.

Conclusion: Physical Design as a Measure of Operational Maturity

Edge computing isn't just about proximity—it's about predictability. Those who engineer for it will define the next decade of distributed intelligence.

The limits to edge success are physical, not computational. Thermal, power, and environmental stability now define uptime as surely as software defines function.

Organizations that treat infrastructure as a strategic design discipline—investing in visibility, modularity, and manufacturability—will outpace those that treat it as an afterthought.

In this new phase, the differentiator isn't whose software runs the workload—it's whose infrastructure keeps it running. The next generation of the edge will be defined not by where it operates, but by how intelligently and reliably it performs. At CPI, we see every enclosure as an opportunity to engineer confidence—turning distributed intelligence into dependable reality.

References

[1, 2] *IDC Estimates Global Spending on Edge Computing to Grow at 13.8% Reaching Nearly $380 Billion by 2028*
https://my.idc.com/getdoc.jsp?containerId=prUS53261225

[3, 4, 5] *5G and AI Convergence at the Edge*
https://www.wipro.com/engineering/5g/navigating-the-convergence-of-5g-and-ai-at-the-edge/

[5] *The Convergence of AI and IoT in 2025*
https://smartdev.com/the-convergence-of-ai-and-iot-in-2025/

[6] *Optimizing Edge AI with Advanced Thermal Management in Embedded Systems*
https://www.embedded.com/optimizing-edge-ai-with-advanced-thermal-management-in-embedded-systems-2/

[7] *Data Center Containment Market Driving Energy Efficiency and Thermal Management in IT Facilities*
https://www.precedenceresearch.com/data-center-containment-market

[8, 9] *Annual Outage Analysis 2025*
https://uptimeinstitute.com/resources/research-and-reports/annual-outage-analysis-2025

[10] *2025 Trends in Edge Computing Security*
https://www.otava.com/blog/2025-trends-in-edge-computing-security/

Contributors

Michael Moore
Sr. Product Manager, Open Architecture Systems

Michael Moore has more than 25 years of experience within the telecommunications industry. He has held various Product Management roles where he was responsible for developing, launching, and managing new innovative products within the Data Center, Airflow Management, and Information Technology market segments.

David Knapp
Product Marketing Manager

David Knapp has more than 20 years of experience in the telecommunications industry with CPI as a product-application expert and technical communicator in the roles of Technical Support, Technical Writer and Product Marketing Manager. He is currently focusing on data center, enterprise networking, including industrial networks, and power management solutions. David enjoys identifying, developing and sharing innovative ways of protecting your technology investment.

chatsworth.com
techsupport@chatsworth.com
800-834-4969

CHATSWORTH PRODUCTS