WHITE PAPER

CHATSWORTH PRODUCTS

Scaling Infrastructure at Speed: How Communication, Capacity, and Engineering Alignment Enable Global Expansion

By
Todd Schneider,
Director of Product Management, CPI



Published: May 2026

chatsworth.com

Executive Summary



Infrastructure growth has entered a new phase where the constraint is less about whether technology exists and more about whether organizations can execute consistently at scale. Synergy Research Group's market analysis shows hyperscale operators continue to rapidly expand and are taking a growing share of global data center capacity[1]. At the same time, McKinsey's analysis highlights accelerating demand for AI-ready data center capacity through 2030, driven by the power density and infrastructure requirements of advanced AI workloads[2].

In this environment, infrastructure outcomes increasingly depend on a partnership model built on four reinforcing pillars:

Hyperscale Partnership Model Pillars			
Continuous Communication	**Shared Solution Development**	**Capacity Planning Beyond the Data Center**	**Intentional Tradeoffs**
Early and ongoing alignment between end user engineering, operations, and manufacturing counterparts to manage change without field disruption.	Infrastructure platforms collaboratively designed to support evolving structural, electrical, thermal, and liquid cooling requirements—helping organizations scale more efficiently as density demands increase.	Throughput planning that includes manufacturing capacity, regional production, and supply chain resilience—not just megawatts and square footage.	A decision framework that optimizes **total deployment economics,** not unit cost, by quantifying the downstream impact of quality variance and lead time volatility.

Critically, at mass scale, infrastructure decisions are not owned by one function. They are a **group decision** across **procurement, engineering, and operations**—because each discipline governs a different source of risk: commercial, technical, and deployment lifecycle.

This paper outlines how organizations scaling infrastructure rapidly can operationalize this model to increase rollout speed, reduce replication risk, and sustain predictable global expansion as AI drives higher density, higher power variability, and more complex cooling architectures.

Scaling Infrastructure Has Become an Integrated Operating Model



Today, this operating model is often executed through pre-engineered, modular, and pod-based deployments, where structural, electrical, and thermal systems are validated to work together prior to rollout.

Scaling Infrastructure in the AI Era

Historically, "hyperscale" has been associated with a small group of large cloud operators. Hyperscale can be more accurately defined as the ability to scale infrastructure rapidly, predictably, and globally in response to demand.

Today, however, many organizations are adopting similar operating principles as AI workloads accelerate infrastructure growth and increase deployment complexity. This model is now being adopted beyond traditional hyperscalers, including:

- **NeoCloud providers building GPU-as-a-service platforms**
- **AI factories supporting large-scale model training and inference**
- **Enterprise and colocation operators deploying high-density AI infrastructure**



What unifies these environments is not ownership, but execution requirements: high deployment velocity, increasing density, evolving power and cooling architectures, and the need for repeatable global rollout.

As a result, the model outlined in this paper applies broadly to any organization operating at—or preparing for—this level of scale.

While hyperscalers are advancing this model through fully integrated, reference-based deployments, colocation providers and enterprise operators are more often navigating phased upgrade paths—extending existing environments to support higher density and hybrid cooling requirements.

Large infrastructure deployments used to behave like a portfolio of discrete builds: a facility, a design handoff, a procurement event, and a deployment cycle. That model breaks down when the operating reality becomes continuous expansion—multiple sites in flight, multiple regions ramping, and design standards iterating in parallel with workload evolution.

As expansion velocity increases, the system becomes sensitive to **variance.**

A small dimensional tolerance issue in a cabinet rail doesn't remain a local nuisance; it becomes a global installation disruption when replicated across thousands of units and multiple deployment partners. A modest lead time slip doesn't "delay a shipment"; it interrupts commissioning sequences and causes labor inefficiencies across multiple sites. A minor change in grounding practice or busway interface becomes a repeatable field risk when documentation and training lag behind design updates.

This is why large-scale infrastructure deployment must be treated as an operating discipline—an execution system that integrates:

- **Design and engineering validation**
- **Manufacturing readiness and throughput**
- **Deployment repeatability**
- **Lifecycle serviceability**

That integration cannot sit solely in procurement, nor solely in engineering. Procurement manages supplier risk, commercial alignment, and long-term capacity commitments. Engineering validates the platform against projected density growth and evolving cooling architectures. Operations is accountable for repeatable deployment, efficient serviceability, and sustained lifecycle performance.

When these functions operate in silos, the deployment system accumulates friction. When they act as a coordinated governance body, infrastructure becomes a controllable lever for expansion speed and reliability.



Open, Continuous Communication as a Foundation for Predictability



The cadence of AI-driven infrastructure expansion now exceeds traditional infrastructure cycles. AI training clusters compress schedules and increase density targets. Regional power constraints force architectural adjustments and on-site power generation considerations. Cooling strategies evolve from air-dominant to hybrid to liquid integration. Meanwhile, compliance and localization requirements introduce complexity across global regions.

In this environment, "communication" is not a soft value. It is a risk control.

The most expensive failures are not typically dramatic design flaws. They are **misalignments discovered late,** after manufacturing begins or after shipments arrive on site:

- **Late design changes that trigger rework or hold inventory**
- **Installation interference between components that were never validated together**
- **Site teams improvising because standards evolved but documentation did not**
- **Field retrofits that introduce inconsistency across regions**

Uptime Institute's resiliency research emphasizes that human and process factors contribute to a significant majority of downtime incidents over time, reinforcing how procedural alignment and execution discipline reduce risk[3]. In hyperscale deployments, late change management and inconsistent procedures are exactly the conditions that amplify that risk.

Continuous, cross-functional dialogue—between engineering teams, deployment leaders, manufacturer engineering, and manufacturing operations—creates three practical advantages:

Earlier Validation	Faster Change Management	Higher Confidence in Rollout
Potential conflicts (structural loading, airflow constraints, power interfaces, service clearances) are identified before they become field issues.	When standards evolve, the update path is coordinated across design documentation, BOMs, manufacturing work instructions, and deployment training.	When manufacturing and deployment feedback loops are active, the system adapts without introducing regional "drift."

Once communication becomes continuous, organizations can plan for change instead of reacting to it.

Shared Solution Development

AI has introduced a density inflection point that is changing infrastructure requirements at the cabinet level. McKinsey's work on AI data center growth highlights how AI-ready capacity demand is accelerating and how AI workloads continue to drive higher power density requirements.[2]

What matters operationally is that higher density is no longer a niche scenario. It is becoming a planning baseline across many new builds and expansions, requiring infrastructure platforms that can evolve alongside rapidly changing compute, power, and cooling demands.

As organizations scale, infrastructure providers and operators must work together to ensure platforms align not only to current deployment requirements, but also to future density roadmaps and operational strategies. This shared solution development approach becomes increasingly important across three areas:

 Structural	 Electrical	 Thermal
Cabinet platforms must support heavier compute loads, denser power distribution, and potential liquid cooling integration while maintaining deployment repeatability and long-term serviceability.	Power architectures must adapt to higher sustained loads, changing redundancy models, increasing monitoring requirements, and evolving regional deployment standards.	Airflow management, containment, and liquid cooling integration must operate as coordinated systems capable of supporting both current and future density requirements.

As AI infrastructure evolves, organizations that align engineering, manufacturing, and deployment planning early can reduce redesign cycles, maintain operational consistency, and accelerate deployment readiness across future expansion programs.

Capacity Planning Beyond the Data Center



Large-scale infrastructure planning has historically emphasized megawatts, land, and timelines. That's no longer sufficient. Manufacturing throughput and supply chain resilience are now first-order constraints—particularly as AI pushes demand for electrical and cooling infrastructure.

McKinsey has noted the scale of investment and the acceleration pressure on data center buildouts, which increases sensitivity to supply constraints[2].

An organization can have land, permits, and power commitments in place—and still miss deployment targets if upstream manufacturing cannot deliver predictably.

Capacity planning at hyperscale now includes:

- **Manufacturing throughput:** Can suppliers produce at the cadence the rollout requires?
- **Regional production:** Can output be aligned to geography to reduce logistics friction, geopolitical, and environmental impacts?
- **Supply chain resilience:** Can the ecosystem absorb disruption without schedule collapse?

Multiple/secondary suppliers designated in advance: Backup sourcing strategies planned before disruption occurs (rather than improvised during a shortage).

Multi-year agreements are a key lever here because they transform capacity from an uncertain variable into a planned asset: reserving throughput, smoothing lead times, and enabling suppliers to invest in tooling and labor stability.

The practical objective is not to eliminate variability, rather, the objective is to prevent variability from becoming systemic—where a regional shortage forces inconsistent substitutions and creates a fragmented global infrastructure standard.

Rethinking Price, Quality, and Lead Time Through Total Deployment Economics

At hyperscale scale, traditional cost-per-unit thinking becomes misleading because unit cost is only a small component of total outcome.



The real cost drivers are:

- **Deployment schedule slips**
- **Field labor variance**
- **Rework and retrofit rates**
- **Commissioning inefficiencies**
- **Lifecycle service disruption**

This is the logic of **total deployment economics.**

The decision framework needs to be explicit:

- **Price** is meaningful, but not in isolation. It must be evaluated against the downstream cost of variance.
- **Quality** must be evaluated as repeatability at scale. A small defect rate replicated across thousands of units is operationally material.
- **Lead time** must be evaluated as schedule reliability. The cost of unpredictability often exceeds the cost of speed.



A simple replication illustration shows why this matters: If a deployment program rolls out 10,000 cabinets and experiences a 1% issue rate that requires on-site remediation, that's 100 interventions. If each intervention consumes 3–6 labor hours plus scheduling overhead (access coordination, safety protocols, escalation workflows), the labor and schedule impact becomes measurable—especially when those interventions are distributed across regions and teams.

This is why strong partnerships matter: they enable organizations to make **intentional tradeoffs** rather than reactive compromises. When engineering, procurement, and manufacturing are aligned, tradeoffs are planned, quantified, and consistent across the global program.

A Partnership Model for Scaling at Speed

The four pillars discussed so far are not independent tactics. They reinforce each other:

Communication enables shared innovation because requirements are visible early and changes are managed coherently.

Shared solution development enables infrastructure platforms to evolve alongside changing density, power, and cooling requirements.

Long-term planning enables capacity alignment because throughput can be reserved and standardized.

Capacity alignment enables smarter tradeoffs because decisions are based on system economics, not short-term procurement pressure.



This is the partnership model increasingly used by organizations scaling infrastructure rapidly: a coordinated execution framework that treats infrastructure as a repeatable system rather than a sequence of purchases.

What organizations operating at hyperscale should expect from modern manufacturing partners:

1. **Engineering collaboration that begins early and persists through rollout**
2. **Capacity transparency that supports multi-year deployment cadence**
3. **Technology readiness aligned to density and cooling evolution**
4. **Execution discipline that preserves consistency across regions**

Applying the Model: What to Look for in Scalable Infrastructure Platforms



When the partnership model is operationalized effectively, its impact is reflected directly in the architecture and repeatability of the infrastructure platform.

Structural readiness for densification

Cabinet and rack platforms should be engineered for high-load environments with predictable repeatability. That includes structural integrity, accessory compatibility under load, and service clearances that remain consistent when cabinets are populated with dense equipment.

Structural readiness also includes foresight: the ability to add higher density hardware later without requiring a cabinet platform redesign.

Power architecture designed for scale and change

Power distribution must scale across regions and adapt to evolving redundancy strategies. The platform should support:

- **Scalable configurations without redesign**
- **Monitoring granularity aligned to load balancing and operational visibility needs**
- **Compatibility with busway and top-of-cab strategies common in hyperscale**
- **Global adaptability that reduces regional requalification effort**

Airflow management engineered for predictability

At higher densities, airflow becomes less forgiving. Platforms need engineered airflow pathways, effective isolation, and containment strategies that can support mixed density within an aisle. The goal is not only cooling performance—it is **repeatable performance** across sites and installers.

Liquid readiness without locking into a single cooling future

The most practical requirement today is optionality: the ability to integrate liquid cooling selectively without forcing an entire program into one architecture. Liquid-ready design is not just about "can it be done"; it is about whether it can be done consistently, serviceably, and safely at scale.

Deployment consistency across regions

Finally, deployment consistency is where hyperscale programs succeed or stall.

Platforms should support:

- **Global compliance readiness**
- **Standardized installation methodology**
- **Deployment repeatability across teams and geographies**
- **Universal solutions that minimize regional variants and keep the global standard coherent**

System-level integration matters more than individual components because integrated platforms reduce interface risk—fewer surprises in the field, fewer compatibility assumptions, and fewer ad-hoc adaptations.



How CPI Supports Scalable Infrastructure Deployment

CPI supports hyperscale execution by aligning engineering collaboration, manufacturing discipline, and integrated infrastructure platforms to the realities of density evolution and global deployment cadence.

Engineering Collaboration and Pre-Deployment Validation

Early technical engagement enables customer teams to validate structural loading assumptions, airflow modeling, and power distribution architecture before scaled production begins. This reduces late-stage redesign risk and helps ensure that cabinet platforms, power systems, and containment strategies are validated against forward-looking density targets — not only current requirements.

By integrating engineering feedback into manufacturing workflows early, CPI helps minimize field variability and maintain configuration consistency across regions.

Platforms Designed for Repeatable Deployment

CPI's infrastructure platforms are engineered not as standalone components, but as coordinated systems designed for repeatable global rollout.

ZetaFrame® Cabinet System

The ZetaFrame® Cabinet System is engineered is engineered for high-density environments with high load capacity and structural headroom to support evolving compute profiles. Its preintegration readiness supports streamlined staging, reduces field assembly time, and improves deployment repeatability across multiple regions. By accommodating density escalation and liquid integration without structural redesign, the ZetaFrame Cabinet System helps customers avoid mid-cycle cabinet platform shifts.



ORV3 Racks

CPI supports hyperscale and AI-driven environments with USA manufactured ORV3 rack solutions aligned to Open Compute Project (OCP) specifications. Designed for compatibility with OCP architectures, these racks enable seamless integration into standardized deployment models while supporting evolving infrastructure requirements.



CPI's ORV3 capabilities include:

- **Flexible design options:** allowing adaptation to specific deployment standards and facility constraints
- **OCP-compatible mounting:** designed for seamless integration of standardized components, including power shelves
- **Busbar provisioning support:** aligned with OCP power distribution strategies and future configuration flexibility
- **Scalable manufacturing and delivery:** enabling high-volume production aligned to hyperscale rollout timelines

By combining OCP alignment with customization and manufacturing scale, CPI enables organizations to deploy ORV3-based infrastructure with **consistency, flexibility, and speed at scale.**



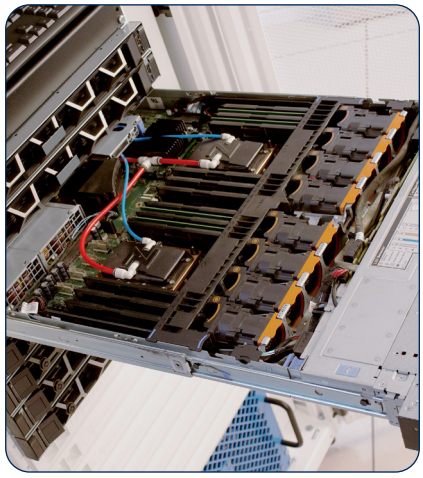
eConnect® PDUs

eConnect® PDUs provide intelligent, scalable power distribution aligned to next-generation load architectures. Branch-level visibility and configurable input options support evolving redundancy models and regional deployment requirements. Including universal PDU solutions reduces configuration fragmentation across global programs and simplifies inventory management at scale.

Aisle Containment

Aisle Containment Build to Spec (BTS) Kits deliver engineered airflow separation and thermal isolation tailored for hyperscale environments. Designed to support variable density and mixed thermal profiles, BTS containment enhances cooling predictability by minimizing bypass and improving airflow control without imposing rigid infrastructure constraints. Its adaptable design helps maintain consistent thermal performance as workloads evolve across sites.





Direct-to-Chip Liquid Cooling Integration (ZutaCore® HyperCool®)

CPI's integrated direct-to-chip liquid cooling approach combines the ZetaFrame® Cabinet System with ZutaCore® HyperCool®, a two-phase, waterless cooling solution designed to support high-density AI and HPC environments. Engineered for cabinet-level integration, the solution supports AI silicon exceeding 2,500W without throttling while helping reduce energy consumption and water usage compared to traditional cooling approaches.

Cabinet-level liquid integration supports a range of densification strategies while preserving operational flexibility. In many hyperscale deployments, liquid cooling is implemented within fully integrated, pre-validated pod or modular architectures, while more incremental, cabinet-level approaches are more common in colocation and enterprise environments.

Coordinated Systems, Not Isolated Components

The objective is not to position any individual product as the solution. It is to support execution through coordinated structural, electrical, and thermal systems that:

- **Reduce integration friction**
- **Improve deployment repeatability**
- **Maintain global configuration consistency**
- **Align infrastructure evolution with long-term expansion programs**

In large-scale environments, execution stability is the differentiator. Integrated platform design and manufacturing alignment help reduce the operational variance that slows global rollout.

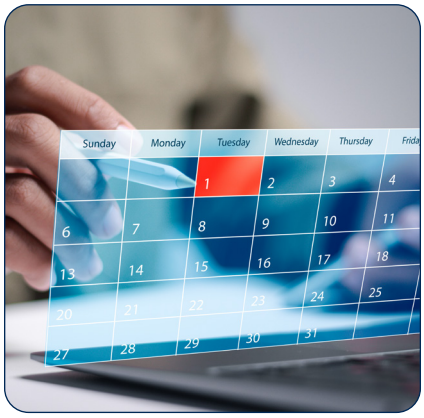


Putting the Partnership Model for Scaling at Speed into Practice

For many organizations, the challenge is not understanding the principles behind large scale infrastructure execution—it is translating those principles into operational behavior across engineering, procurement, and deployment teams.

A practical starting point is evaluating how well your current infrastructure program aligns with the four pillars outlined in this paper.

1. Evaluate Communication Cadence Across Stakeholders



Large-scale execution depends on continuous coordination between design, manufacturing/suppliers, and deployment teams. Ask:

- **Are engineering, procurement, and operations aligned before infrastructure standards change?**
- **Do suppliers receive visibility into future density and cooling roadmap assumptions?**
- **Are field deployment insights consistently fed back into design standards?**
- **If communication happens only at procurement events or late in the deployment cycle, execution risk increases significantly as infrastructure programs scale globally.**

2. Assess Platform Readiness for Density Evolution



AI workloads are pushing cabinet-level infrastructure beyond traditional enterprise assumptions.

Organizations should evaluate whether their infrastructure platforms are prepared for:

- **Higher structural load profiles**
- **Evolving power distribution architectures**
- **Hybrid air-and-liquid cooling integration**
- **Repeatable installation across regions**

Platforms engineered with structural, electrical, and thermal headroom reduce the need for mid-cycle redesign as density requirements evolve.



3. Align Infrastructure Planning with Manufacturing Throughput

Many infrastructure programs plan around facility readiness—power availability, land, and permits—while underestimating the role of manufacturing capacity and supply chain resilience.

Key questions include:

- **Can manufacturing output support the planned deployment cadence?**
- **Are production capabilities aligned regionally with rollout locations?**
- **Are capacity commitments secured through multi-year supplier alignment?**

Treating manufacturing throughput as part of infrastructure capacity planning helps prevent supply constraints from slowing global expansion.

4. Reframe Procurement Around Total Deployment Economics

At scale, infrastructure outcomes are influenced less by unit cost and more by repeatability and schedule reliability.

Procurement teams should evaluate infrastructure decisions based on:

- **Installation labor variability**
- **Deployment schedule stability**
- **Field remediation rates**
- **Lifecycle serviceability**

When these factors are considered alongside price organizations make more consistent infrastructure decisions that support large-scale rollout programs.



A Cross-Functional Alignment Opportunity



Organizations that successfully scale infrastructure treat deployment readiness as a **cross-functional discipline,** not a series of independent procurement decisions.

Evaluating the four pillars together helps identify where execution friction may emerge before the next wave of expansion begins.

Scaling Without Slowing Down

AI-driven infrastructure growth will continue to accelerate—driven by cloud expansion and AI workloads that increase density, power variability, and cooling complexity. The limiting factor will not be whether technology exists. It will be whether organizations can **execute consistently at scale.**

The partnership model offers a practical framework for doing so:

- **Continuous communication that reduces late program risk**
- **Shared solution development helping organizations scale more efficiently as density demands increase**
- **Capacity planning that includes manufacturing and supply chain resilience**
- **Intentional tradeoffs based on total deployment economics**

Organizations that institutionalize this model increase rollout speed, reduce replication risk, and build long-term resilience across global infrastructure programs.



References

[1] Synergy Research Group, *The World's Total Data Center Capacity is Shifting Rapidly to Hyperscale Operators, 2025.*

[2] McKinsey & Company, *AI data center growth: Meeting the demand and related analysis on AI-driven infrastructure expansion, 2024–2025.*

[3] Uptime Institute, *Annual Outage Analysis Reports, 2024–2025.*

Contributors



Todd Schneider,
Director of Product Management, CPI

Todd Schneider is the Director of Product Management and Hyperscale Business Development at Chatsworth Products, Inc. Todd has over 25 years experience in the data center space covering a wide variety of functional areas including Product Management, Business Development, Sales Engineering and Product Development. His current roles focus on driving standard new product innovation to CPI's core offerings and delivering customer-centric, customized solutions for Hyperscale Operators. Schneider previously worked at Legrand, Electrorack and Wright Line.

About Chatsworth Products

At Chatsworth Products (CPI), it is our mission to address today's critical IT infrastructure needs with products and services that protect your ever-growing investment in information and communication technology. We act as your business partner and are uniquely prepared to respond to your specific requirements with global availability and rapid product customization that will give you a competitive advantage. At CPI, our passion works for you. With over two decades of engineering innovative IT physical layer solutions for the Fortune 500 and multinational corporations, CPI can respond to your business requirements with unequaled application expertise, customer service and technical support, as well as a global network of industry-leading distributors. Headquartered in the United States, CPI operates from multiple sites worldwide, including offices in Mexico, Canada, China, the United Arab Emirates and the United Kingdom. CPI's manufacturing facilities are located in the United States, Asia and Europe.

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

CHATSWORTH PRODUCTS

While every effort has been made to ensure the accuracy of all information, CPI does not accept liability for any errors or omissions and reserves the right to change information and descriptions of listed services and products.

©2026 Chatsworth Products, Inc. All rights reserved. Chatsworth Products, Clik-Nut, CPI, CPI Passive Cooling, CUBE-iT, Secure Array, eConnect, Evolution, GlobalFrame, MegaFrame, QuadraRack, RMR, Saf-T-Grip, SeismicFrame, SlimFrame, TeraFrame, Motive, Velocity, Oberon, Wi-Tile and ZetaFrame are federally registered trademarks of Chatsworth Products. EuroFrame, H-Plane, Hi-Bar, In-Plane, M-Frame,NetPoint, Simply Efficient, Skybar, Elevate, and VersaEdge Wall-Mount Cabinet are trademarks of Chatsworth Products and Oberon, a division of Chatsworth Products. All other trademarks belong to their respective companies. 07/26 MKT-CPI-869