



RESEARCH BRIEF

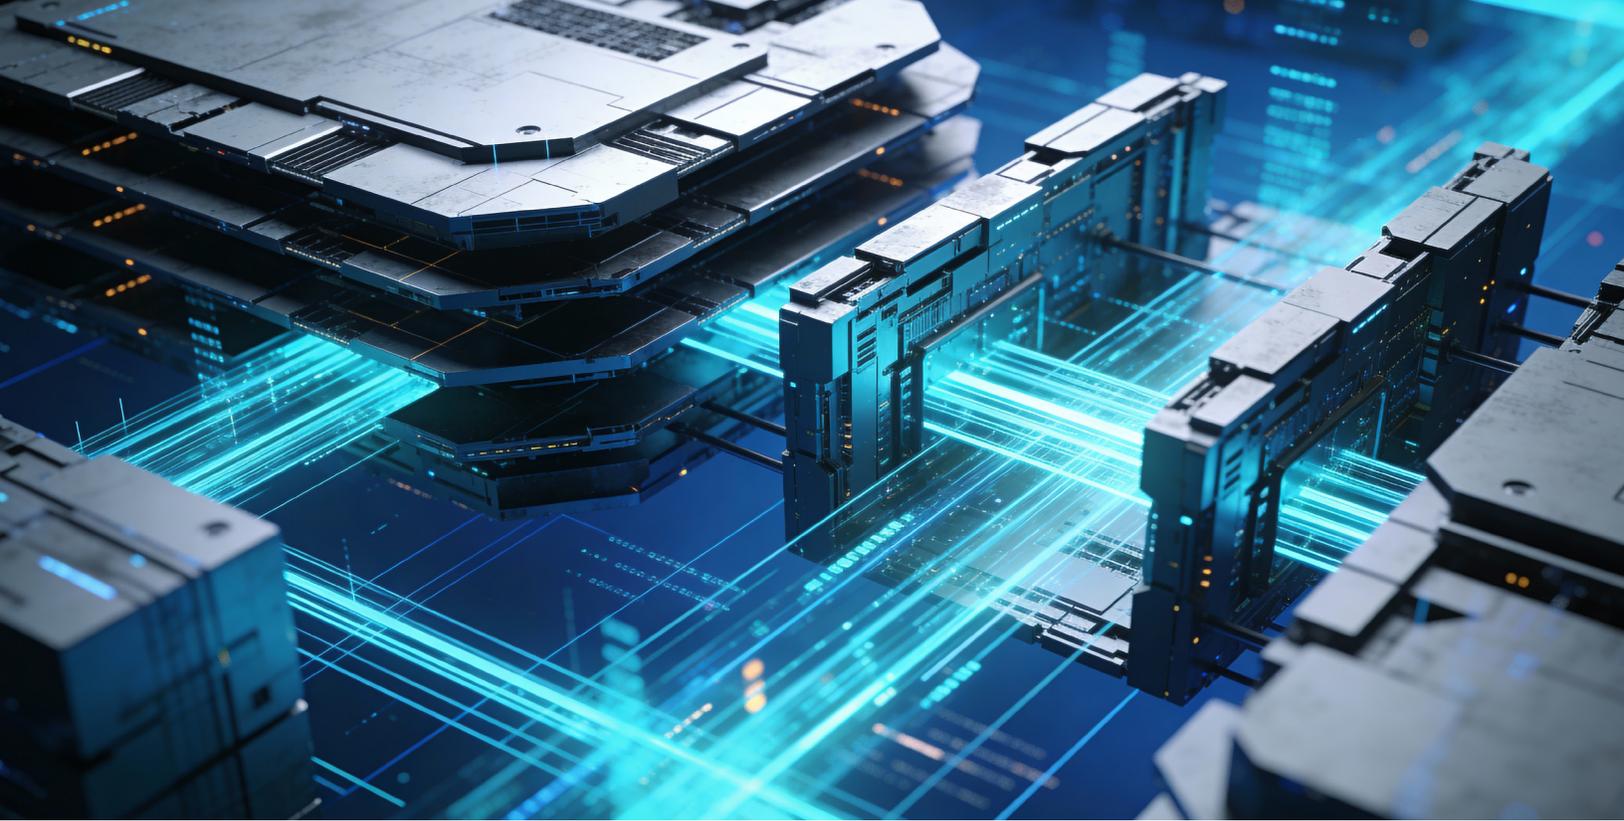
# The Hyperspeed Compute Era Reclaiming AI Velocity for Enterprise Teams

Authors:

**Stephen Sopko**  
Analyst-in-Residence –  
Semiconductors & Deep Tech

**Steven Dickens**  
CEO & Principal Analyst

MARCH 2026



## Executive Summary

Enterprise AI development is at an inflection point. The infrastructure layer, no longer an afterthought or vanilla part of the operations plan, has become the primary choke point for competitive velocity. This leaves organizations desiring the fail-fast-iterate approach to AI development with a significant “compute crunch” characterized by high scarcity-driven GPU/hardware costs and uncertainty about availability of the infrastructure needed for AI training and inference. This can leave even the largest firms struggling to secure necessary resources.

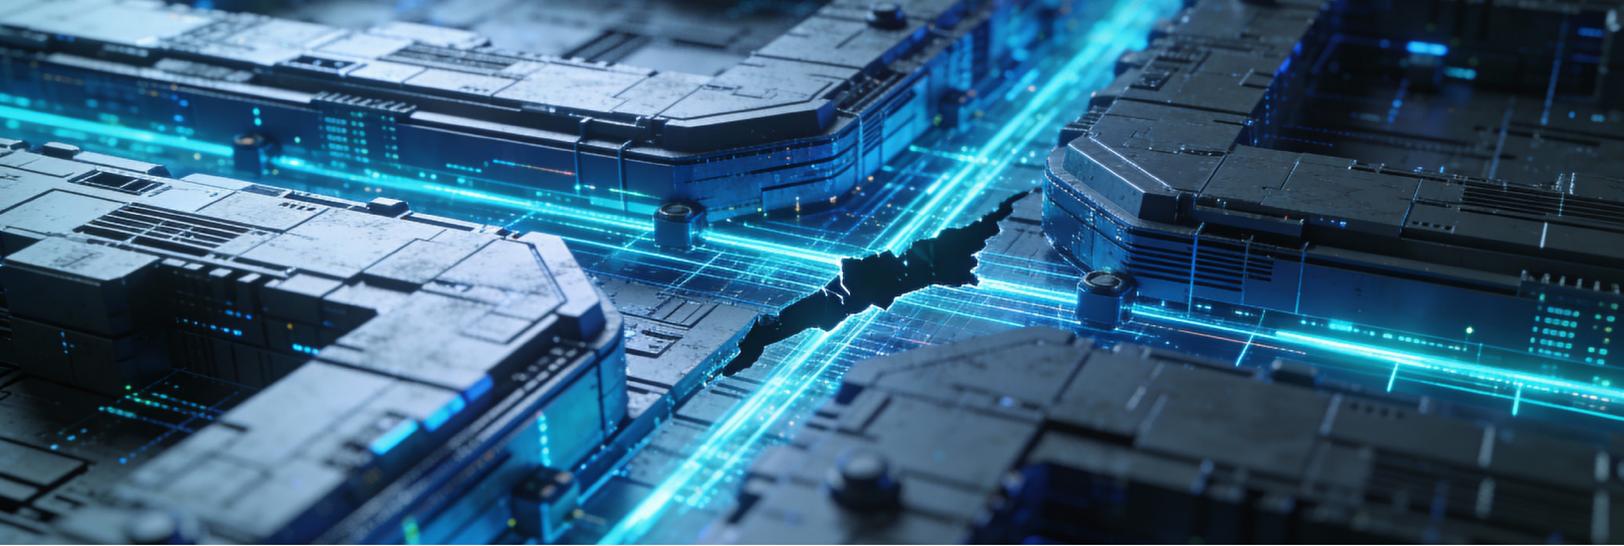
Compounding the physical bottlenecks are severe shortages of the skilled talent who can effectively manage complex AI data pipelines. Finally, with all of these hurdles, organizations in top-tier cities encounter the energy scarcity needed to power data centers. Consequently, the trajectory of AI adoption in the coming age of inference is shifting away from experimentation and giant pilots into a “flight to efficiency.” In this movement, companies increasingly select the option of smaller, fine-tuned open-source models over the massive, custom-built infrastructures. This creates a growing divide between “AI-mature” organizations that can navigate these overheads and those stuck in “pilot purgatory” due to mounting technical debt. Ultimately, the future of enterprise AI will be defined by strategic resource orchestration rather than just raw model capability.

Against this backdrop, organizations with detailed velocity plans for AI talent and model development, and billions in investment waiting on the sidelines, often find themselves:

- Waiting weeks for GPU capacity,
- Losing crucial burn rate on idle commitments,
- Navigating pricing structures optimized for different workload patterns.

This Research Brief examines why yesterday’s infrastructure model built around information-scale workloads cannot support the development of offerings and solutions at intelligence-scale. This Research Brief will also introduce the QumulusAI FACTS framework (Flexibility, Access, Cost, Trust, Speed) as a diagnostic and decision lens for infrastructure decision-makers, and present QumulusAI’s approach to “hyperspeed compute”: a distributed architecture designed to eliminate the queues, commitments, and opacity that characterize legacy GPU provisioning.

**Key Finding:** Organizations establishing infrastructure velocity advantages in early 2026 create the opportunity for those advantages to compound across every subsequent AI development cycle. Infrastructure choice is no longer an operational hurdle; it is the strategic enabler.



# 1. The Infrastructure Velocity Gap

## 1.1 When Infrastructure Became the Bottleneck

Global AI infrastructure spending is projected to accelerate its historic upward curve, with annual growth rates estimated by various analysts between 20%-40% through 2029. This surge reflects a fundamental shift: now infrastructure points like power and facilities are the critical path of AI development instead of model development and training. A football-stadium sized hyperscaler data center takes years to design, permit, provision, and deploy. Organizations that can bypass the long lines in the AI supply and infrastructure chains to provision, iterate, and scale fastest are going to be the organizations that survive and thrive.

Yet the dominant infrastructure model was architected to meet different constraints: breadth of service and economies of scale serving diverse enterprise workloads. This hyperscale approach of massive centralized data centers serving virtualized workloads across global networks was built when the bottleneck was storage and networking, not processor speed and access to infrastructure.

## 1.2 The Workload Architecture Alignment Challenge

General-purpose cloud architectures and AI-native workloads present intrinsically different optimization requirements. Hyperscale infrastructure is perfect for workloads that define information-scale: serving web pages, processing transactions, streaming content, storing documents. They can eventually serve intelligence-scale workloads (training models, running inference, fine-tuning on proprietary data) but struggle with fundamentally different challenges at intelligence-scale:

- **Burst capacity requirements:** AI development cycles oscillate between dormant windows and intensive training runs requiring maximum GPU allocation. Commitment-based pricing forces teams to pay for capacity they don't use or wait for capacity they need. This creates tension between infrastructure economics and AI development patterns.
- **Provisioning latency:** Multi-week wait times for GPU allocation are common in hyperscale environments. For teams built around fail-fast iterations of model architectures, every week of waiting extends iteration cycles and delays learning.
- **Cost opacity:** Layered pricing for storage, access, bandwidth, and compute creates budgeting uncertainty that discourages intelligent experimentation. This drives teams to optimize for conservative cost predictability rather than needed performance, with the resulting loss of crucial time to value.
- **Ecosystem lock-in:** Proprietary orchestration tools and managed services are seen as moats against competitors, creating prohibitive switching costs that compound over time. Organizations must avoid yet another generation of tech-debt by weighing integration benefits versus long-term flexibility requirements and vendor behaviors.

These aren't failures of the hyperscaler model. They reflect architectural priorities optimized for breadth of service at global scale. Organizations pursuing AI-native workloads benefit from evaluating whether general-purpose infrastructure aligns with their specific development patterns.

### 1.3 The Cost of Waiting

Infrastructure friction symptoms present differently depending on the organizational context. Two illustrative scenarios reveal the patterns:

#### Scenario A: The AI Developer's MVP-to-Production Journey

In this scenario, an AI development firm is building private/sovereign LLM products, including customized small language models, for their customers. Their business model requires speed and margin flexibility: they only survive by demonstrating product-market fit with minimized capital outlay, then scale capacity as customer demand develops. Under the legacy infrastructure, this team faces a fundamental innovation roadblock. The small-scale experimentation needed for fail-fast is expensive relative to results. Legacy-scale commitments require capital they cannot deploy until they achieve demonstrable product-market fit. The resulting cash flow squeeze and development impediment delays time-to-market and constrains the experimentation necessary to accelerate product development.

For these teams, the challenge of expensive GPU time is a transient issue, driven by current scarcity that should reduce over time. The real question for this scenario is how quickly can the team move from proof-of-concept to production without betting the company on capacity commitments and deployment delays?

#### Scenario B: The Enterprise Internal AI Team

Now consider an enterprise with multiple business units pursuing independent and competing AI initiatives. Each unit needs capacity for experimentation, but none can justify dedicated infrastructure at scale until use cases are validated. Under centralized procurement models, honed to work with the legacy hyperscale model, teams queue for shared resources in the cloud. Internal and vendor SLAs create multi-week delays and massive uncertainty that kills projects. Budget cycles prohibit the rapid scaling of capacity needed for breakthrough value-creation. The only groups that succeed use shadow IT (its back!) that can fragment vendor commitment load that could drive down cost, and AI initiatives lose momentum waiting for infrastructure that never arrives. Additionally, eventual infrastructure consolidation can force teams to restart development processes, compounding earlier delays.

The common thread: infrastructure friction not only slows development and kills fail-fast, it shapes which projects get approved or attempted. Teams stop proposing experiments they know will get stuck in provisioning queues, and shift to either AI-absent solutions, or back channels that fragment organizational leverage with vendors. The opportunity cost is staggering, but invisible until competitors achieve their breakthroughs.



## 2. The FACTS Framework

Every organization driving to implement AI at intelligence-scale confronts the same infrastructure challenges. The result is a sense that everybody is waiting in the same line for capacity. The FACTS framework from QumulusAI seeks to provide a diagnostic and decision lens, and allows for infrastructure choices that empower development velocity. QumulusAI designed its architecture around these five dimensions.

### QumulusAI FACTS Framework

AI Infrastructure for Intelligence-Scale



### F – Flexibility

**The Challenge:** AI workloads are inherently variable: training runs demanding burst capacity, inference requires sustained throughput at scale. Experimentation in both needs to be low-commitment as value is quantified and proven. Existing hyperscale infrastructure and business models drive teams to choose between intentional over-provisioning (thus paying for idle capacity) and more conservative under-provisioning (resulting in a wait for availability). Fixed-tier models can inhibit the rapid experimentation AI development requires.

**The Hyperspeed Approach:** QumulusAI's architecture enables seamless transitions across deployment models. Serving up fractional GPU access for prototyping all the way through

dedicated instances for scaling to bare-metal clusters for production workloads. This flexible approach gives organizations the ability to right-size infrastructure to current workflow, rather than committing to legacy-model capacity tiers designed for steady-state operations.

This flexibility extends to scale. QumulusAI delivers it with deployments under 50MW distributed across colocation partnerships. This lets customers bypass the capacity concentration and resulting allocation bottlenecks in mega-scale data centers. Another key part of flexibility is a portfolio approach to AI development. With hyperscaler assets used as appropriate, and hyperspeed where the need for rapid deployment or a tailored solution is indicated.

## A – Access

**The Challenge:** Provisioning delays are difficult to plan for, but can kill AI development momentum. Multi-week wait times for GPU allocation, and the reduced ability to iterate on development can compound into multi-month project delays. Teams may lose institutional knowledge and momentum as engineers and subject-matter-experts context-switch to other projects while waiting for hyperscale capacity.

**QumulusAI's Hyperspeed Approach:** QumulusAI's distributed architecture is achieved by constantly replenishing GPU capacity across colocation partnerships. This can eliminate the centralized infrastructure allocation queues with provisioning measured in hours, not weeks. Customer development cycles can now iterate at the pace of learning rather than the pace of procurement bureaucracy.

For enterprise teams, this transforms capacity planning from a quarterly exercise into an operational variable. Scale up when experiments succeed and adoption takes off; scale down when priorities shift or lessons learned must be implemented. The infrastructure dynamically adapts to the business rather than constraining it.

## C – Cost

**The Challenge:** GPU infrastructure pricing is the strategic variable, not just an operational cost line item. The wild expansion of the GPU and memory markets in the age of AI often defeat the transparency and predictability needed for even medium-term implementation strategy. The existing model features complex pricing structures with egress fees, variable

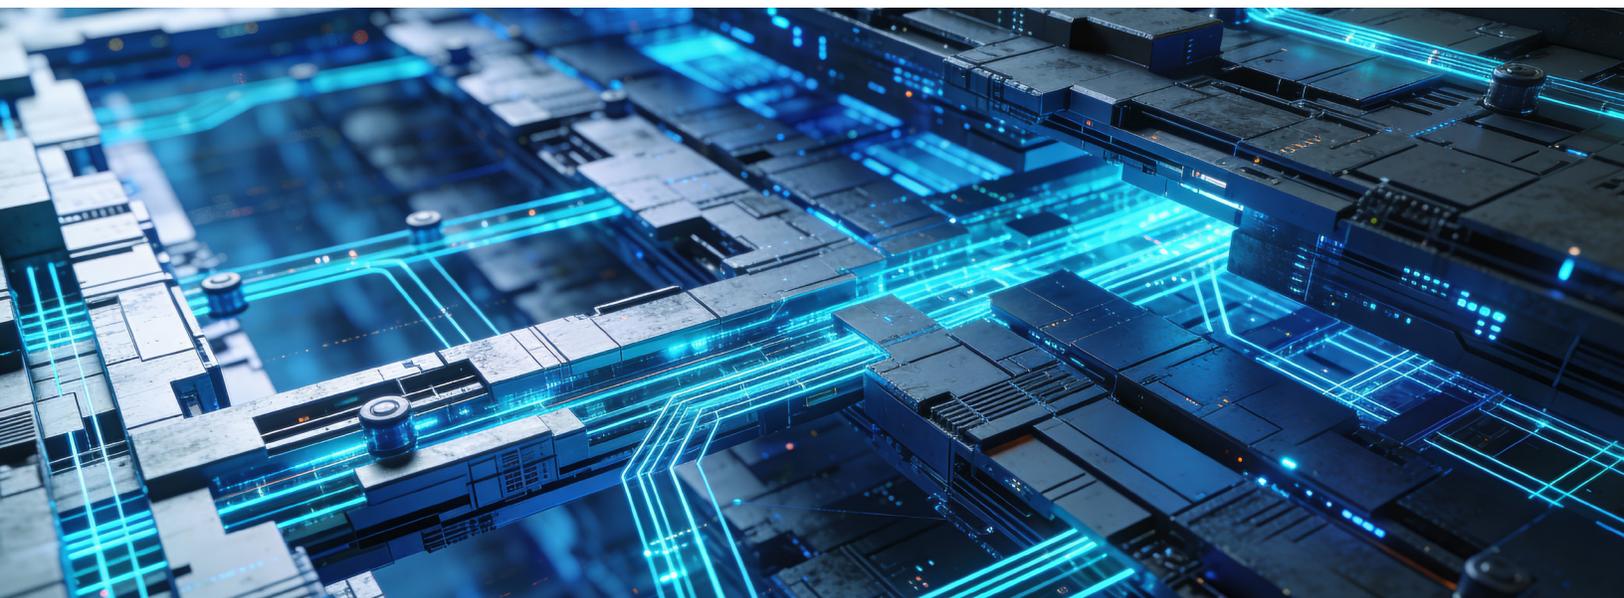
storage charges, and commitment penalties; all of which reduce the efficacy of cost forecasting that then discourages downstream experimentation and innovation.

**QumulusAI's Hyperspeed Approach:** QumulusAI's business model seeks to deliver better pricing transparency. No hidden egress fees. No unpredictable storage charges. The QumulusAI approach to pricing structures looks to enable teams to forecast monthly costs with confidence. This transparency aims to deliver sustained experimentation and the kind of iterative development that builds into competitive advantage.

More importantly, the QumulusAI relationship model is designed to shift the relationship from commodity-vendor to innovation-partner. Capacity planning as a service is a concept where enterprises work with partners to anticipate requirements rather than react to constraints after the fact. The infrastructure relationship becomes a strategic partnership rather than a procurement transaction.

## T – Trust

**The Challenge:** Enterprise infrastructure decisions carry multiple risk profiles highly dependent on the organization and its regulatory framework. Downtime impacts production systems. Security vulnerabilities can expose proprietary data. SLA failures cascade into customer commitments. These risks are not addressed well by the typical one-size-fits-all offerings often advocated for by the hyperscaler providers, leaving risk on the enterprise without the flexibility to design for specific needs. Decision-makers need confidence that infrastructure partners can deliver consistently.



**QumulusAI's Hyperspeed Approach:** Trust is built through visibility and relationship, but also through a flexible approach. QumulusAI is looking to deliver clear usage metrics, proactive capacity forecasting, and direct engineering access so that customers can own more of the control to match the risk based on their specific needs. Enterprise teams work with dedicated contacts to build the trust infrastructure necessary for their unique workloads and priorities.

This operational transparency, matched with flexibility, makes clear not just capacity levels, but the unique processes and protections that enterprise AI initiatives require.

## S – Speed

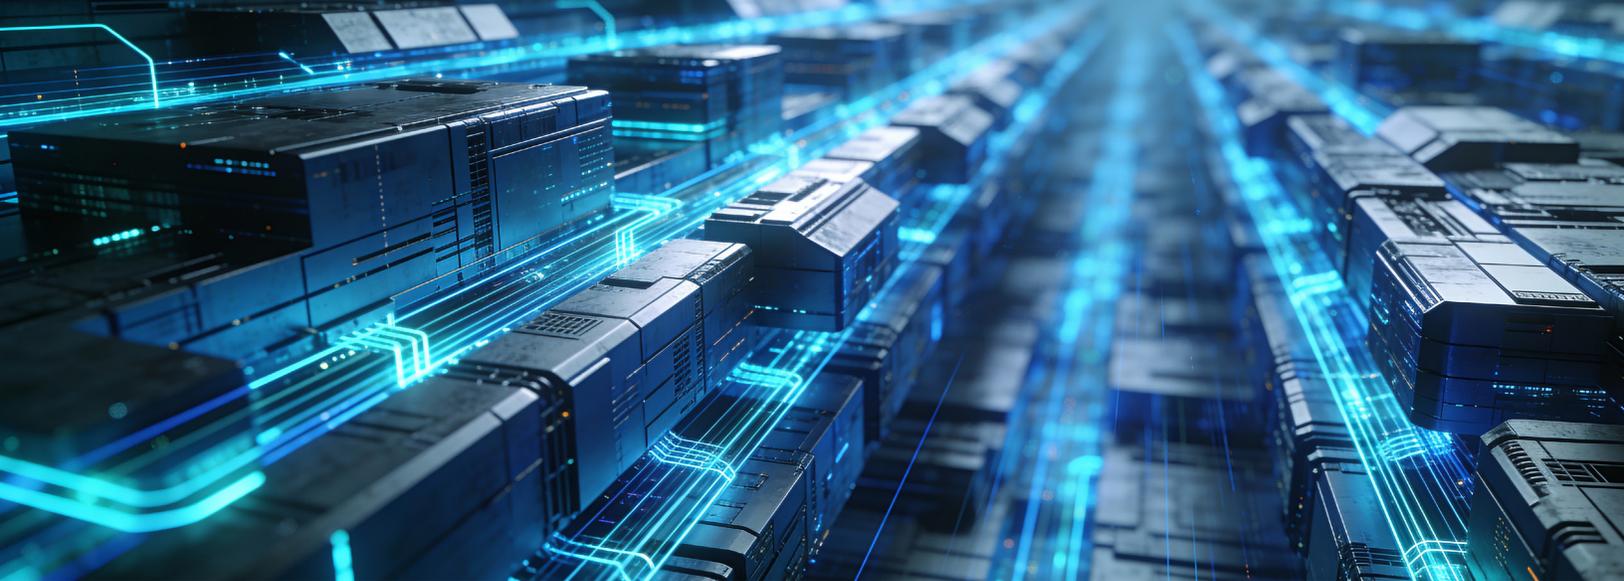
The Challenge: Traditional information-scale workloads are generally more tolerant of delay or latency than intelligence-scale workloads. A quarterly reporting cycle doesn't necessarily suffer if provisioning takes time, because it is a predictable load factor. AI development operates under dramatically different timescales than most enterprise applications: model architectures are constantly evolving, training runs generate insights requiring immediate follow-up, and AI-powered products must go to market rapidly before the market evolves past the opportunity. Speed is multi-layered in AI development:

- Provisioning speed determines how quickly teams can start.
- Iteration speed determines how quickly they can learn and implement lessons learned.
- Scaling speed determines how quickly they can capitalize on success.

Bottlenecks at any dimension constrain the entire development cycle. General-purpose cloud infrastructure, optimized for diverse enterprise workloads, may introduce friction at each of these layers for AI-specific development patterns.

**QumulusAI's Hyperspeed Approach:** QumulusAI's architecture sees speed as a first-order design constraint, not a secondary consideration. As a result, the company works with infrastructure partnerships to empower QumulusAI customers to bring on provisioning in a fraction of the time. Seamless scaling as experimentation turns into success at scale. Hyperspeed infrastructure accelerates development cycles rather than acting as a speed-brake constraining them. This is the essence of hyperspeed compute: infrastructure velocity yielding competitive differentiation.





## 3. Rethinking Infrastructure Strategy

### 3.1 The Emerging GPU Cloud Landscape

The AI infrastructure market is fragmented into distinct segments, each with different strengths and constraints:

- Hyperscale providers offer global reach and deep service ecosystems, but at premium pricing and reduced flexibility. Architectures are optimized for different workload patterns, and suffer from deployment lag.
- GPU marketplaces aggregate distributed capacity at attractive price points, but those lower prices are achieved with variable hardware quality, inconsistent availability, unclear SLA's, and limited enterprise support.
- Emerging GPU cloud providers deliver AI-native infrastructure, but with varying approaches to capacity, deployment flexibility, and customer relationships. They use the legacy hyperscale model instead of hyperspeed.

#### Portfolio Approaches to AI Infrastructure

Modern enterprise AI strategies increasingly involve multiple infrastructure partners rather than single-provider commitments. Hyperscaler ecosystems offer compelling value for integrated ML pipelines, global-scale production inference, and workloads benefiting from tight integration with broader cloud services. Specialized GPU cloud partners can complement these deployments by providing rapid-access capacity for experimental workloads, burst requirements, and development phases where iteration speed outweighs ecosystem integration.

The strategic question isn't which provider to choose exclusively, but how to architect the right infrastructure mix for each workload phase and organizational requirement.

**QumulusAI** therefore intends to occupy a distinctive position:

- Infrastructure designed by and for AI development teams,
- With a distributed architecture built around hyperspeed access,
- And the partnership orientation (up and downstream) supporting enterprise planning.

The company's focus is on eliminating the friction that slows development velocity rather than building sprawling service ecosystems.

### 3.2 Infrastructure as Partnership, Not Transaction

The traditional enterprise infrastructure relationship is transactional and commoditized: capacity purchased, consumed, and renewed. This model treats infrastructure as interchangeable, inflexible, and optimized primarily on price for the convenience of corporate procurement.

QumulusAI's approach aims to reframe infrastructure as a partnership. Reducing the risk of infrastructure ownership by seeing capacity planning as a service. **That approach requires: working together to anticipate requirements, designing focused solutions to direct customer needs; thus moving from reactive procurement to strategic planning.** The relationship adds value beyond raw compute, delivering the flexibility needed for innovation and demonstrating the value of AI solutions in real offerings rather than pitch decks.

For enterprise decision-makers, this distinction matters. Transactional vendor relationships optimize for the next procurement cycle. A partner relationship optimizes for the next capability. Working together to identify the capacity and structures needed for innovation. As AI becomes increasingly central to competitive positioning at all levels of the market, the difference between commodity and partnership compounds.

### 3.3 The Pilot Program Advantage

Given the rapid evolution and high stakes of AI-enabled enterprises, infrastructure transitions require hard facts validation without overcommitment. QumulusAI's pilot program model enables enterprises to test the hyperspeed approach with a contained scope and clear evaluation criteria.

Phase

1

#### Assessment:

Use the FACTS dimensions to map current infrastructure friction points. Identify high-impact opportunities to translate velocity gains into real-world opportunities.

Phase

2

#### Validation:

Deploy pilot workloads with defined success metrics. Compare provisioning speed, iteration speed, and cost predictability against baseline infrastructure.

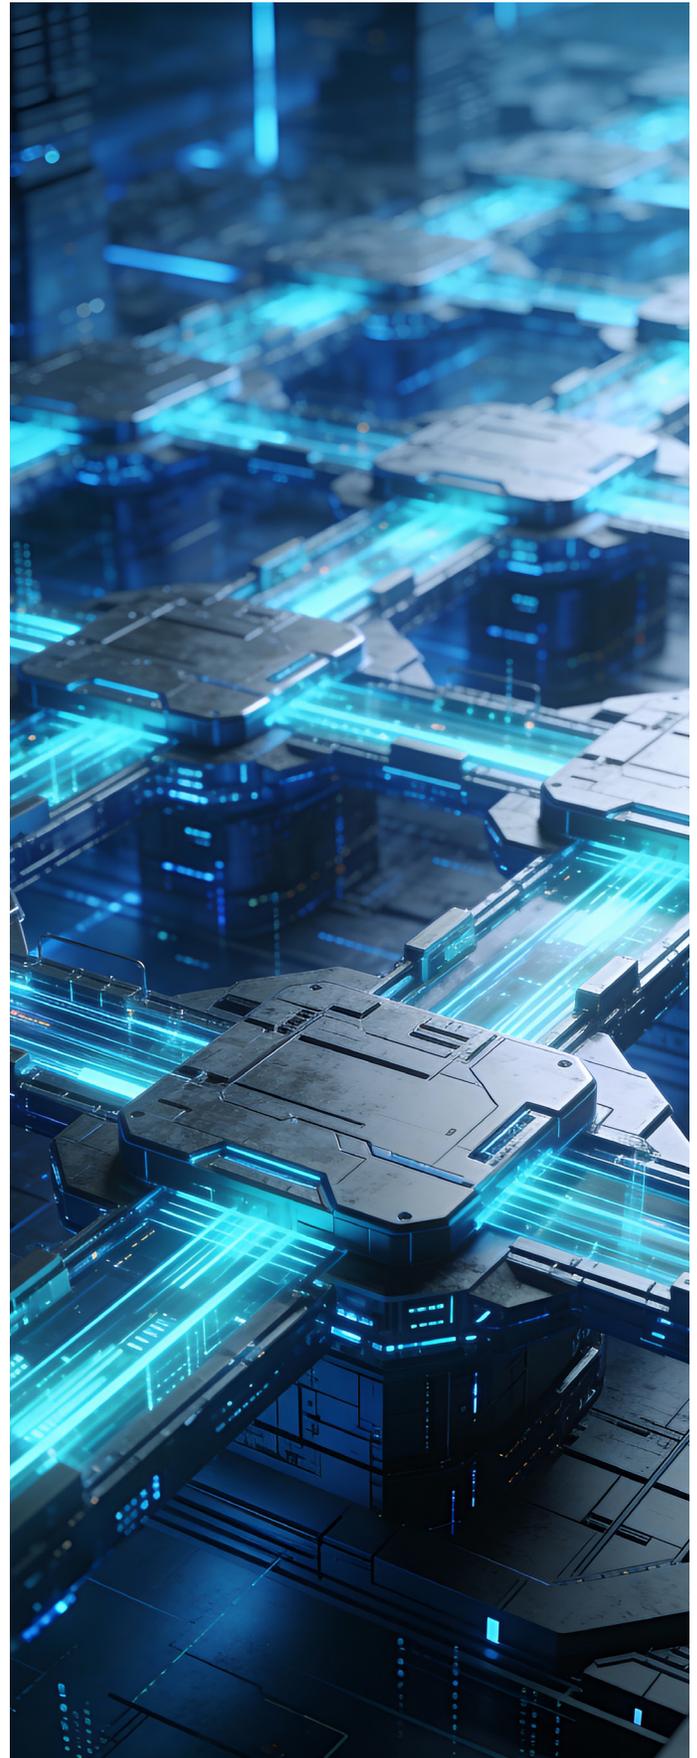
Phase

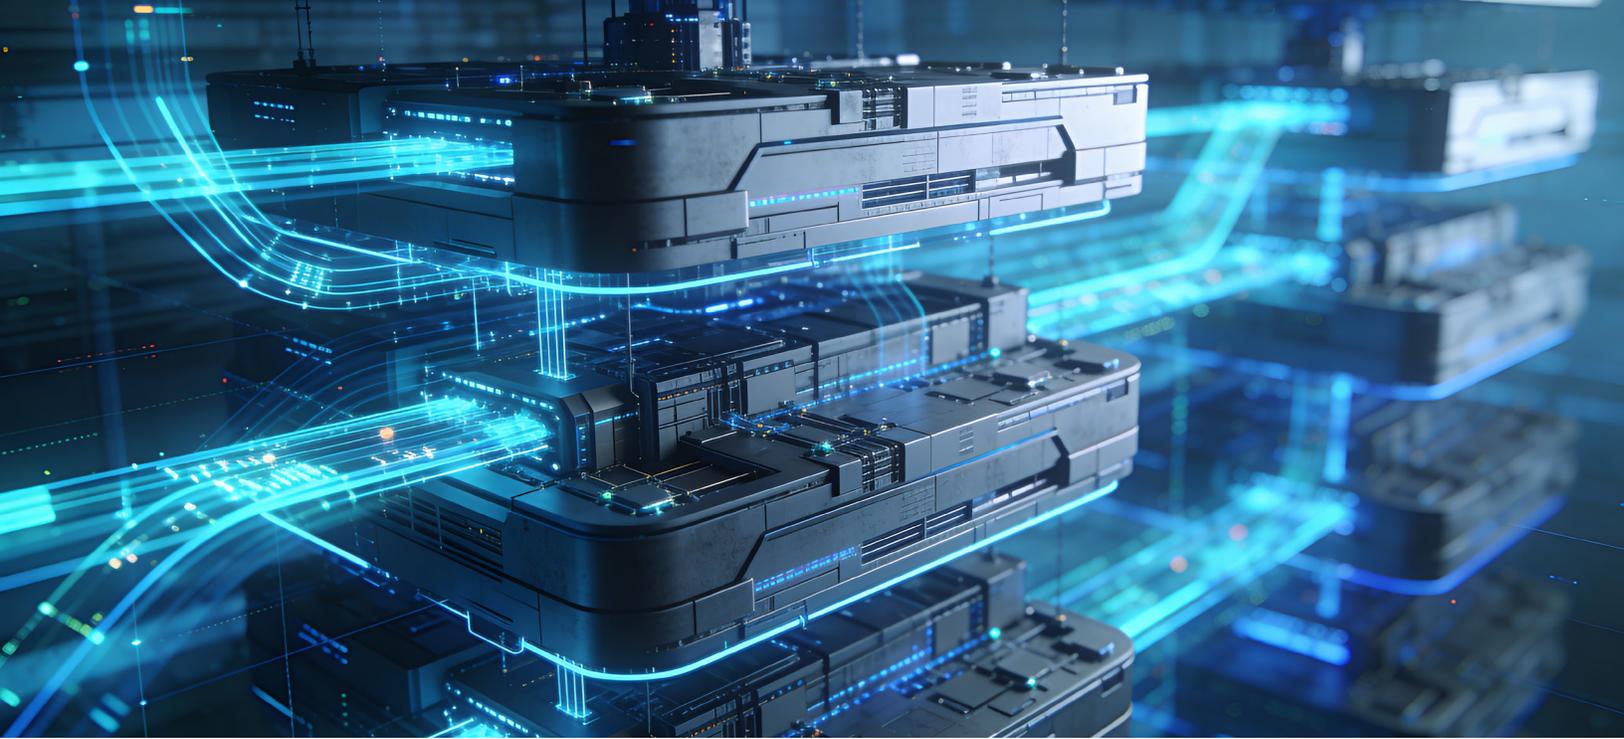
3

#### Scaling:

Where successful pilots demonstrate the right pattern, use the model flexibility and implement the capacity forecasting partnership. Transform infrastructure into a strategic enabler rather than its current position of unpredictable constraint.

Phased piloting is the very definition of rapid experimentation and iteration, giving organizations the ability to validate hyperspeed compute benefits with their own workloads and requirements.





## 4. Building Enterprise AI Strategy

### 4.1 Traditional Infrastructure vs FACTS Hyperspeed

The difference between legacy and hyperspeed infrastructure is the difference between waiting and iterating. It is the difference between infrastructure as a development cycle constraint and infrastructure as an innovation enabler.

Under general-purpose models optimized for diverse enterprise workloads, provisioning delays can create stop-start development patterns. Teams submit requests, wait for allocation while trying to hold on to resources, execute intensive work, learn from the results, then return to waiting and trying to hold on to executive attention and needed developers. Knowledge inevitably decays between cycles. Momentum dissipates. Projects that should take months take years - but more often are abandoned entirely. This contributes to the 'AI doesn't work for us' perception.

Under the QumulusAI FACTS approach, infrastructure can become a continuous flow. Provision as needed. Scale as enterprise requirements grow. Iterate as learning happens, not as procurement and supply chains enable. Shorter projects turn controversial resource-holds from intra-organizational warfare into a non-issue. Continuous development cycles (faster learning, faster iteration, faster deployment) delivers a competitive flywheel of the sustainable competitive advantage at the heart of the AI promise.

### 4.2 Key Questions for C-level Decision-Makers

The QumulusAI FACTS approach looks to enable infrastructure assessment through the lens of these key diagnostics

#### 1. Is provisioning latency and capacity friction killing your AI momentum?

**Key checks:** Have projects been delayed, de-prioritized, or abandoned due to multi-week GPU wait times? Are iteration cycles stuck below 1–2 per month because teams lose momentum/context while waiting? Is infrastructure availability the primary bottleneck to faster learning/experimentation?

**Implication for executives:** If “yes,” your current setup enforces a “wait in line” model incompatible with AI’s fail-fast, iterative nature. This directly erodes competitive velocity in 2026’s intelligence-scale era, where weeks of delay compound into lost market opportunities. A portfolio approach (hyperspeed + hyperscaler) addresses this with pinpoint provisioning and distributed capacity replenishment, turning infrastructure from a brake into an accelerator.

#### 2. Are cost unpredictability and idle commitments draining budgets without delivering value?

**Key checks:** Do hidden/egress fees, variable charges, or commitment lock-ins create monthly budget surprises and discourage experimentation? Are teams over-provisioning (paying for idle capacity in troughs) or under-provisioning (risking delays when experiments succeed)? Are high up-front commitments blocking rapid, low-risk iteration?

**Implication for executives:** If “yes,” infrastructure costs are strategic drag rather than controlled enablers. Legacy models force trade-offs between predictability and performance; FACTS prioritizes transparent, usage-based pricing (no hidden fees) and seamless up/down scaling, enabling confident forecasting and sustained experimentation without burning runway on idle resources.

### 3. Does your infrastructure limit strategic optionality and long-term flexibility?

**Key checks:** Are switching costs, ecosystem lock-in, or rigid scaling constraining which AI projects get pursued? Can you quickly scale up for validated successes or down when priorities shift without penalties/restarts? Does the relationship feel transactional (commodity vendor) rather than partnership-oriented for capacity planning?

**Implication for executives:** If “yes,” infrastructure is capping your AI portfolio and creating invisible opportunity costs. A portfolio approach, blending hyperscalers for steady-state/production with a solution like QumulusAI for burst/experimental velocity, restores optionality. QumulusAI’s

model emphasizes partnership (proactive forecasting, direct engineering access) and flexibility (fractional to bare-metal, distributed <50MW deployments), positioning infrastructure as a strategic multiplier rather than a constraint.

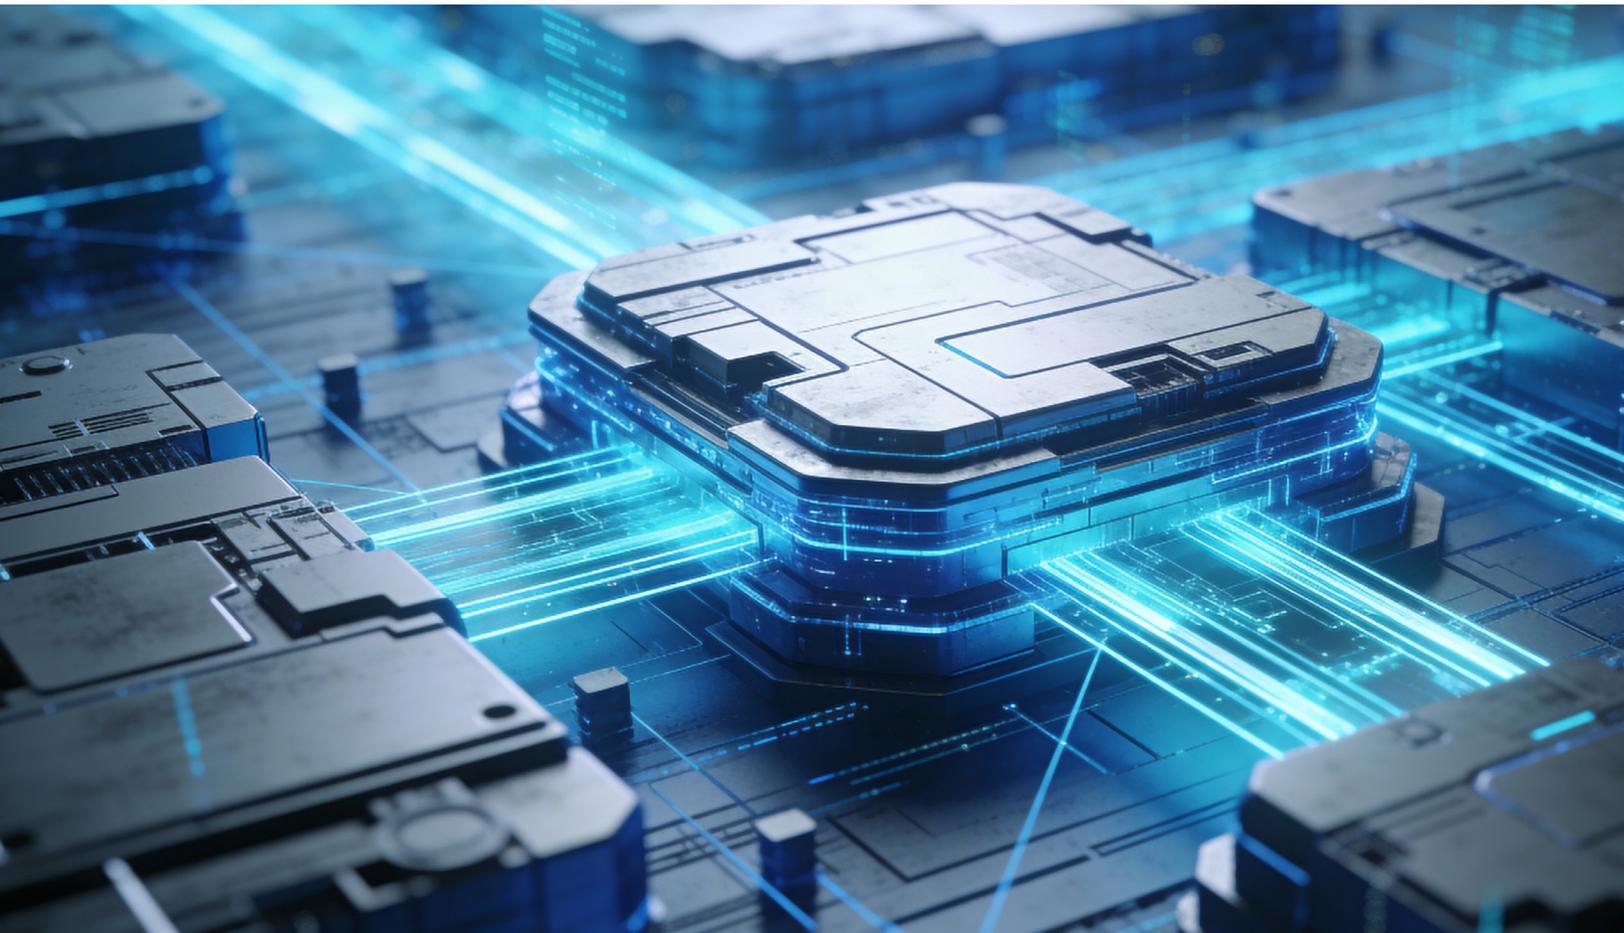
### 4.3 Implementation Roadmap

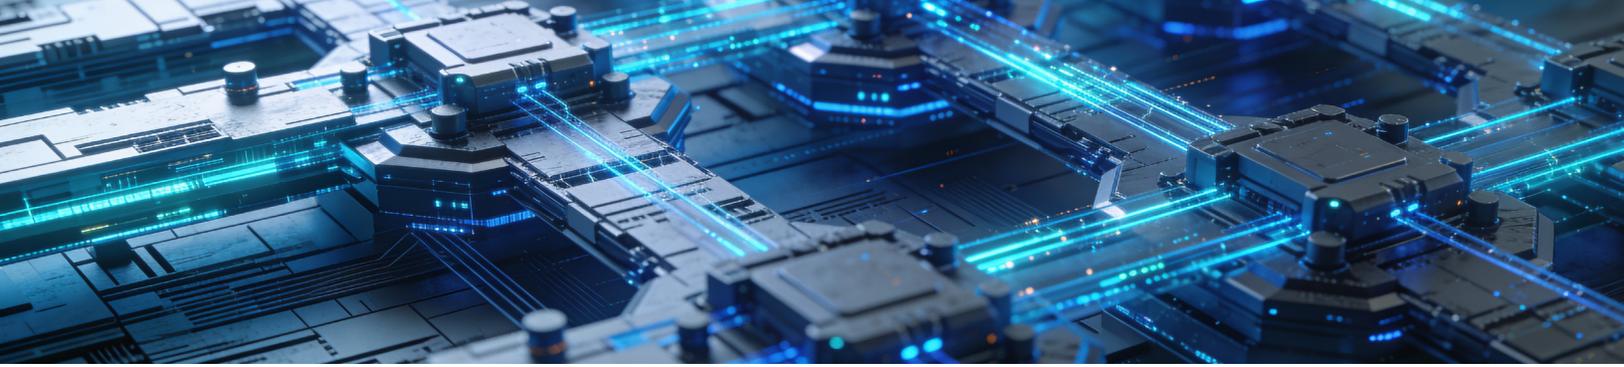
Organizations transitioning to hyperspeed compute typically follow this three-phase approach:

**Assessment:** Audit current infrastructure friction using FACTS dimensions. Quantify delay costs and opportunity costs. Identify pilot workloads that can show clear benefits from higher velocity.

**Validation:** Deploy targeted workloads with clear success metrics. Validate the speed of provisioning, cost transparency, and flexibility of up/down scaling. Build a joint business case targeting wider adoption based on lessons learned.

**Strategic Scaling:** Expand the case to additional teams and workloads. Build this expansion into the capacity planning partnership. Transition infrastructure availability beyond constraint to competitive advantage.





## 5. The Hyperspeed Imperative

### 5.1 Infrastructure Flex Is Strategic

AI development flexibility is no longer a second-order requirement; it is now central to success. Organizations that deploy faster will iterate faster and learn faster. These organizations compound advantages across every subsequent development cycle. Infrastructure choice and flexibility is the foundation of this compounding.

The hyperscale model is optimized (and demonstrably works) for breadth, for global reach, for service ecosystems at global providers. These remain valuable for AI, but don't address the AI bottleneck. The AI bottleneck is still GPU provisioning, and will be for the foreseeable future. The bottleneck results in reduced development velocity, and frustrates the ability to move from experimentation to production without waiting for infrastructure. Opening the bottleneck reduces delays and increases transparency - both essential for innovation.

### 5.2 The Window for Competitive Advantage

Infrastructure decisions made in the next 12 months establish a foundation for the crucial next three years of AI development. Organizations securing an early velocity advantage benefit from a compounding effect across every subsequent cycle. Organizations that remain locked into legacy hyperscale infrastructure models will be able to try fewer projects, and will find the gap widening.

So the window is about when to start the compounding effect, not just deploying a technology. Every week of provisioning delay is a week of delayed learning, or a week of burned runway to prove results. Every constrained or abandoned experiment is an unexplored optionality. The infrastructure choices enterprises make today determine the AI advantage against competitors tomorrow.

### 5.3 Looking Ahead

As enterprises strive to deploy AI successfully and at scale, they face the challenge of accessing the right infrastructure with the necessary flexibility and pricepoint to ensure the ROI of their project. While numerous options exist for savvy enterprises, two things are paramount in making an infrastructure provider choice: modularity and speed of deployment. Based on HyperFRAME Research's analysis of the market for AI infrastructure provision, the path to hyperspeed compute begins with understanding an enterprise's current friction points and evaluating alternatives against a rubric similar to QumulusAI's FACTS framework.

**Assess enterprise infrastructure friction:** Use the diagnostic questions to quantify where the current infrastructure constrains development velocity.

**Explore a pilot engagement:** QumulusAI's pilot program enables validation with a contained scope and clear success metrics. Test hyperspeed business models by matching workload to the right level of compute-power sooner rather than later based on external and geopolitical forces.

**Develop a strategic infrastructure roadmap:** Go from reactive procurement towards proactive capacity roadmapping. Build infrastructure relationships that accelerate rather than constrain enterprise AI development trajectory. Develop a blended portfolio of hyperscalar and hyperspeed solutions.

Put simply, organizations that establish infrastructure velocity advantages in 2026 will define the competitive landscape of the AI era. The hyperspeed model provides an alternative to 'waiting in line' with everybody else. Carving out a competitive advantage when it is still early enough to matter.



This research brief was prepared by HyperFRAME Research in collaboration with QumulusAI. HyperFRAME Research provides independent analysis of technology infrastructure markets. For questions about this research or to explore QumulusAI's pilot program, contact: [info@qumulusai.com](mailto:info@qumulusai.com)

#### ABOUT HYPERFRAME RESEARCH:

HyperFRAME Research delivers in-depth research and insights across the global technology landscape, spanning everything from hyperscale public cloud to the mainframe and everything in between. We offer strategic advisory services, custom research reports, tailored consulting engagements, digital events, go to market planning, message testing, and lead generation programs.

Our industry analysts specialize in rigorous qualitative and quantitative assessments of technology solutions, business challenges, market forces, and end user demands across industry sectors. HyperFRAME Research collaborates closely with your Analyst Relations, Product, and Marketing teams to build and amplify your thought leadership, positioning your expertise to enhance brand and product recognition. Through content that engages readers, viewers, and listeners alike, we ensure your voice resonates across channels.

#### CONTACT HYPERFRAME RESEARCH:

**Steven Dickens**

CEO & Principal Analyst | HyperFRAME Research

**Email Address:**

[steven.dickens@hyperframeresearch.com](mailto:steven.dickens@hyperframeresearch.com)

**Telephone Number:**

+1 845 505 1678

**X:** [@StevenDickens3](#)

**LinkedIn:** Steven Dickens

**BlueSky:** Steven Dickens

#### CONTRIBUTORS

**Stephen Sopko**

Analyst-in-Residence  
Semiconductors & Deep Tech

**Steven Dickens**

CEO & Principal Analyst  
HyperFRAME Research

#### INQUIRIES

Contact us if you would like to discuss this report and HyperFRAME Research will respond promptly.

#### CITATIONS

This paper can be cited by accredited press and analysts, but must be cited in-context, displaying author's name, author's title, and "HyperFRAME Research." Non-press and non-analysts must receive prior written permission by HyperFRAME Research for any citations.

#### LICENSING

This document, including any supporting materials, is owned by HyperFRAME Research. This publication may not be reproduced, distributed, or shared in any form without the prior written permission of HyperFRAME Research.

#### DISCLOSURES

HyperFRAME Research provides research, analysis, advising, and consulting to many high-tech companies, including those mentioned in this paper. No employees at the firm hold any equity positions with any companies cited in this document.

