

QumulusAI Introduces "Hyperspeed Compute" as a New Model for Enterprise AI Infrastructure

2026-03-05

New HyperFRAME research finds infrastructure velocity is now the primary constraint on enterprise AI progress

ATLANTA, GA / [ACCESS Newswire](#) / March 5, 2026 / QumulusAI, a distributed AI infrastructure provider, today announced the release of a new research brief developed in collaboration with **HyperFRAME Research**, *The Hyperspeed Compute Era: Reclaiming AI Velocity for Enterprise Teams*. The report examines why enterprise AI initiatives are increasingly stalled by infrastructure constraints and outlines a new approach designed to eliminate lengthy GPU access delays, rigid capacity commitments, and cost opacity.

According to the research, enterprise AI has entered a "flight to efficiency" phase. Rather than large, monolithic model builds, teams are prioritizing smaller, fine-tuned models and faster iteration cycles. Further, most infrastructure environments remain optimized for "information-scale" workloads - pages, processing transactions, streaming content, storing documents - instead of "intelligence-scale" workloads - training models, running inference, fine-tuning on proprietary data. The result is a widening infrastructure velocity gap that separates AI-mature organizations from those stuck in prolonged pilots.

"The biggest limiter on enterprise AI today isn't models or ambition, it's access," said **Mike Maniscalco, CEO of QumulusAI**. "Teams are waiting weeks, if not months, for GPU capacity, paying for idle commitments, and losing momentum while procurement and provisioning catch up. Infrastructure has become a strategic bottleneck - and teams should be looking to augment hyperscale infrastructure with hyperspeed compute."

Infrastructure Is Now the Competitive Choke Point

The HyperFRAME report identifies three structural issues shaping enterprise AI outcomes in 2026:

- **Provisioning latency:** Multi-week waits for GPU access slow iteration and kill a fail-fast development strategy.

- **Architectural misalignment:** Hyperscale environments optimized for steady workloads struggle with burst-driven AI development.
- **Cost uncertainty:** Complex pricing models and commitment structures discourage experimentation.

These constraints not only slow projects, they shape which AI initiatives are attempted at all.

"Infrastructure choice now directly determines AI velocity," said **Steven Dickens, CEO and Principal Analyst at HyperFRAME Research**. "Organizations that remove friction early gain a compounding advantage across every development cycle that follows."

Introducing Hyperspeed Compute and the FACTS Framework

The research introduces QumulusAI's **FACTS** framework - **Flexibility, Access, Cost, Trust, and Speed** - as a diagnostic lens for evaluating AI infrastructure readiness. The framework is designed to help enterprises identify where legacy infrastructures create friction and where alternative architectures can restore development momentum.

QumulusAI's approach, described in the report as "**hyperspeed compute**," is built around:

- Flexible scaling from fractional GPUs to dedicated clusters
- Access to distributed GPU capacity across colocation partners
- Cost transparency without hidden egress or storage fees
- Trust-based partnership model focused on capacity planning, not one-time transactions
- Speed - rapid deployments designed to bring compute online for clients in weeks vs months

The report recommends a **portfolio approach**, combining hyperscale environments for steady-state workloads with hyperspeed infrastructure for experimentation, burst capacity, and early production phases.

From Waiting to Iterating

The report concludes that infrastructure decisions made in early 2026 will shape enterprise AI competitiveness for years to come. Organizations that prioritize infrastructure velocity can iterate faster, learn faster, and deploy faster - creating a flywheel effect that compounds over time.

The full research brief, *The Hyperspeed Compute Era*, is available from QumulusAI (ADD LINK). Enterprises interested in validating the approach can participate in QumulusAI's pilot program, designed to test provisioning speed, cost predictability, and iteration velocity with real workloads.

About HyperFRAME Research

HyperFRAME Research provides independent analysis of AI, cloud, and infrastructure markets, helping enterprises and technology providers understand emerging architectures and their business impact.

About QumulusAI

QumulusAI is a vertically integrated AI infrastructure company focused on delivering a distributed AI cloud by innovating around power, data center and GPU-based cloud services-the company delivers immediate access to high-performance computing with enhanced cost control, reliability, and flexibility. Machine learning teams, AI startups, research institutions, and growing enterprises can now scale their AI training and inference workloads quickly and cost effectively. For more information, visit <https://www.qumulusai.com>

For more information on QumulusAI:

Press: media@qumulusai.com

Investors: investors@qumulusai.com

Follow QumulusAI on social media: <https://www.linkedin.com/company/qumulusai>

This press release contains certain "forward-looking statements" that are based on current expectations, forecasts and assumptions that involve risks and uncertainties, and on information available to QumulusAI as of the date hereof. QumulusAI's actual results could differ materially from those stated or implied herein, due to risks and uncertainties associated with its business and leadership changes. Forward-looking statements include statements regarding QumulusAI's expectations, beliefs, intentions or strategies regarding the future, and can be identified by forward-looking words such as "anticipate," "believe," "could," "continue," "estimate," "expect," "intend," "may," "should," "will" and "would" or words of similar import. Forward-looking statements include, without limitation, statements regarding future operating and financial results, QumulusAI's plans, objectives, expectations and intentions, and other statements that are not historical facts. QumulusAI expressly disclaims any obligation or undertaking to disseminate any updates or revisions to any forward-looking statement contained in this press release to reflect any change in QumulusAI's expectations with regard thereto or any change in events, conditions or circumstances on which any such statement is based in respect of its business, the strategic partnership or otherwise.

SOURCE: QumulusAI

View the original [press release](#) on ACCESS Newswire