NEWS RELEASE

QumulusAI and Shadeform Deploy Two NVIDIA H200 Clusters Totaling 680 GPUs for Leading AI Inference Platforms

2026-05-28

Two-year contract supports AI infrastructure demand from two high-growth AI inference platforms, marking a significant commercial milestone in QumulusAI's hyperdistributed GPU build-out

ATLANTA, GA / ACCESS Newswire / May 28, 2026 / QumulusAI, a provider of GPU-powered cloud infrastructure for artificial intelligence, and Shadeform, a unified GPU platform connecting AI teams with compute resources, today announced a two-year contract for the deployment of two NVIDIA H200 clusters - a 61-node and a 24-node configuration - totaling 85 nodes at QumulusAI's Kansas City location. The deployment supports compute demand from two high-growth AI inference platforms, including one of the fastest-scaling inference networks currently operating at production scale.

The partnership combines QumulusAI's Hyperspeed Deployment capabilities and hyperdistributed data center footprint with Shadeform's ability to connect and match enterprise demand to best-fit GPU infrastructure - delivering scalable, enterprise-grade H200 environments to AI platforms operating at production scale.

"This deployment reflects exactly how we want to grow - through strategic partnerships that tie long-term, high-quality demand to our expanding infrastructure footprint," said Mike Maniscalco, CEO, QumulusAI. "Shadeform's ability to bring committed, multi-year clients to the table accelerates our mission to transform how enterprises of any size can rapidly obtain access to the GPU-powered cloud infrastructure they need for production AI."

"QumulusAI's ability to deploy at hyperspeed and hold committed capacity is exactly what our customers need, said Ed Goode, CEO, Shadeform. "This deployment gives AI inference platforms on Shadeform access to dedicated, enterprise-grade AI infrastructure as they scale their businesses."

For QumulusAI, the contract validates its capacity-forward model - purpose-built to match enterprise-grade infrastructure with committed, long-duration demand. The Kansas City deployment is supported by a $45 million convertible note facility from ATW Partners, announced separately, with $15 million funded to date.

For Shadeform, the partnership demonstrates its ability to deliver large-scale, dedicated GPU capacity at competitive economics -- directly addressing the reliability and cost challenges AI platforms face when scaling inference workloads.

The Kansas City facility is part of QumulusAI's hyperdistributed network of GPU clusters spanning U.S. colocation and company-owned data centers, with more than 150 MW of available capacity and sub-90-day deployment cycles for fully operational GPUaaS environments.

About Shadeform

Shadeform is a unified GPU platform that enables AI teams to discover, compare and deploy compute from every major cloud. With a vetted network of over 30 cloud partners, Shadeform helps AI teams move faster, spend smarter, and scale larger. Learn more at shadeform.com

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

Disclaimer: 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. 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. 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, partnerships or otherwise.

SOURCE: QumulusAI