



NEWS RELEASE

QumulusAI Signs \$240.9 Million in New AI Inference Customer Contracts for NVIDIA Blackwell B300 Capacity

2026-09-29

The agreements cover more than 2,000 B300 GPUs and expand relationships first announced on July 28 and July 23, 2026, taking the two customers' combined contracted commitments with QumulusAI past \$344 million. Extension options would bring the new agreements' total potential value to \$401.6 million over five years

ATLANTA--(BUSINESS WIRE)-- QumulusAI (Nasdaq: QMLS), a neocloud infrastructure provider purpose-built for the AI computing era, today announced two three-year take or pay agreements with AI inference customers. The agreements have combined contracted value of \$240.9 million and cover more than 2,000 NVIDIA Blackwell B300 GPUs across two dedicated clusters.

\$211.9 million, three years: an AI inference platform that helps companies deploy large language models, vision models, speech models and other AI applications with low latency and high reliability. This is the customer's second agreement with QumulusAI. It follows the \$71.9 million, three-year agreement announced July 28, 2026, and brings the customer's contracted commitments with the company to more than \$283 million. Capacity is targeted for turnover to the customer in early 2027.

\$29.1 million, three years: a generative-AI inference platform serving developers and enterprises running image, video and other generative models. This is the customer's second agreement with QumulusAI. It follows the \$32 million, two-year agreement announced July 23, 2026, and brings the customer's contracted commitments with the company to more than \$61 million. Capacity is expected to be turned over to the customer in December 2026.

"We're excited to keep growing alongside AI builders that are changing what this industry can do, and it says a lot that both of these customers came back to commit to three more years with us," said Michael Maniscalco, CEO of QumulusAI. "Inference platforms are planning capacity years out now, and they are backing those plans with cash. That lets us place GPUs against signed contracts rather than against forecasts."

Conditions and qualifications. The \$240.9 million of contracted value covers the three-year initial

terms only. The larger agreement carries two one-year extension options valued at \$70.6 million each, which would bring its total potential value to \$353.1 million over five years. The smaller agreement carries a two-year extension option valued at \$19.4 million, which would bring its total potential value to \$48.4 million. The options may or may not be exercised, and their value is excluded from the contracted commitments stated above.

About QumulusAI

QumulusAI is a distributed AI cloud platform that delivers accelerated access to high-performance GPU compute. Through an inference-first, demand-led deployment model across a network of data center sites, QumulusAI brings compute closer to customer demand, helping AI teams and enterprises scale production AI workloads with speed, flexibility and control. By combining rapid deployment with flexible private cloud infrastructure, QumulusAI gives customers a faster, more adaptable path beyond the capacity constraints of traditional centralized and hyperscale cloud models. Learn more at [QumulusAI.com](https://qumulusai.com).

Follow us on LinkedIn (linkedin.com/company/qumulusai) and X @QumulusAI.

Forward-Looking Statements

This press release contains forward-looking statements within the meaning of Section 27A of the Securities Act of 1933, as amended, and Section 21E of the Securities Exchange Act of 1934, as amended. These include statements regarding the company's agreements with two AI inference platform providers for NVIDIA Blackwell B300 capacity; the anticipated value and term of the agreements; the potential exercise of the extension options and the total potential value of the agreements; the customers' total contracted commitments with the company; the expected timing of capacity turnover under each agreement; and the company's continued growth. Words such as "anticipate," "believe," "estimate," "expect," "intend," "can," "may," "plan," "project," "target," "will" and similar expressions are intended to identify forward-looking statements. These statements are based on management's current expectations and assumptions as of the date of this release. They are subject to risks and uncertainties that could cause actual results to differ materially, including, among others: the company's dependence on a limited number of large customers; delays in the delivery of GPUs and other equipment the company has on order; the availability and cost of power, network connectivity and specialized hardware such as graphics processing units; the company's substantial capital requirements and access to financing, including equipment lease financing; competition and rapid technological change in the high-performance computing and AI markets; the company's limited operating history and history of net losses; and those described in the "Risk Factors" section of the company's registration statement on Form S-1, as amended (File No. 333-292514), and the company's quarterly report on Form 10-Q for the quarter ended June 30, 2026, as amended, filed with the U.S. Securities and Exchange Commission (SEC), as such factors may be updated in the company's subsequent filings with the SEC. QumulusAI undertakes no obligation to update or revise any forward-looking statement, whether as a result of new information, future developments or otherwise, except as required by applicable law.

View source version on businesswire.com:

<https://www.businesswire.com/news/home/20260929733963/en/>

Investor Contact

investors@qumulusai.com

Media Contact

media@qumulusai.com

Source: QumulusAI