

// INVESTOR PRESENTATION

Breaking AI's Biggest Barriers

AI compute online in months, not years. We use existing capacity to meet demand now, while simultaneously building the pipeline to power what's next.

// DISCLAIMER

Forward–looking Statements and Disclaimers

Presentation

This presentation has been designed to provide general information about QumulusAI, Inc. ("QumulusAI" or the "Company"). Any information contained or referenced herein is suitable only as an introduction to the Company. Except as otherwise indicated, the information contained in this presentation speaks only as of the date set forth on the cover of this presentation, and nothing herein shall create any implication that such information is correct as of any time subsequent to such date.

The information contained in this presentation is for informational purposes only. The information contained herein does not constitute or form a part of, and should not be construed as, any offer for sale or subscription of, or any invitation to offer, buy, or subscribe for, any securities, nor shall there be any offer, solicitation, or sale in any jurisdiction in which such offer, solicitation, or sale would be unlawful.

This document is not a prospectus. The information contained in this presentation is not investment or financial product advice and is not intended to be used as the basis for making an investment decision. Neither the Company nor any of its affiliates makes any representation or warranty, express or implied, as to the fairness, accuracy, completeness, or correctness of any information or opinions contained in this presentation, and no reliance should be placed on them. This presentation has been prepared without taking into account the investment objectives, financial situation, or particular needs of any particular person.

The trademarks included herein are the property of their respective owners and are used for reference purposes only. Such use should not be construed as an endorsement of the platform and solutions of QumulusAI.

Forward-Looking Statements

This presentation contains forward-looking statements within the meaning of Section 27A of the Securities Act of 1933, as amended (the "Securities Act"), and Section 21E of the Securities Exchange Act of 1934, as amended (the "Exchange Act"), that reflect the Company's current expectations and beliefs with respect to, among other things, its financial condition, results of operations, plans, objectives, future performance, and business. When used in this presentation, the words "anticipate," "believe," "could," "estimate," "expect," "guidance," "intend," "may," "on track," "plan," "project," "should," "target," "will," "would," and similar expressions are intended to identify forward-looking statements, although not all forward-looking statements contain such identifying words.

Forward-looking statements include all statements that are not historical facts. These statements are based on information available at the time they are made and on management's good-faith beliefs and assumptions as of that time with respect to future events. Such forward-looking statements are subject to various risks and uncertainties. Accordingly, there are important factors that could cause actual outcomes or results to differ materially from those indicated in these statements, including 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 quarterly period ended June 30, 2026, as amended, filed with the U.S. Securities and Exchange Commission (the "SEC"), as such factors may be updated from time to

time in the Company's subsequent filings with the SEC. These factors include, among others, our dependence on a limited number of large customers; the availability and cost of power, network connectivity, and specialized hardware such as GPUs; our substantial capital requirements and access to financing; competition and rapid technological change in the high-performance computing (HPC) and artificial intelligence (AI) markets; and our limited operating history and history of net losses. An investment in the Company entails a high degree of risk, and there can be no assurance that the Company's objectives will be achieved or that investors will receive a return on their investment.

Forward-looking statements may include statements about the Company's future financial performance, including the Company's expectations regarding net revenue, operating expenses, and its ability to achieve and maintain future profitability; the Company's business plan and ability to effectively manage growth; anticipated trends, growth rates, and challenges in the Company's business, particularly in the fields of High-Performance Computing (HPC) and Artificial Intelligence (AI); the timing, scale, and pace of our GPU deployments and related capacity expansion; further development of the Company's facilities and customer base for related services; beliefs and objectives for future operations; trends in revenue, cost of revenue, and gross margin; trends in operating expenses, including technology and development expenses, sales and marketing expenses, and general and administrative expenses, and expectations regarding these expenses as a percentage of revenue; increased expenses associated with being a public company; and other statements regarding the Company's future operations, financial condition, and prospects.

There is no assurance that any forward-looking statements will materialize. You are cautioned not to place undue reliance on forward-looking statements, which reflect expectations only as of this date. QumulusAI does not undertake any obligation to update or review any forward-looking statement, whether as a result of new information, future developments, or otherwise, except as required by applicable law.

Projections and Estimates

This presentation may contain estimates, projections, and other illustrative financial and operating information, including financial and operating data identified with an "E" suffix, which reflect Company estimates. Such information is presented for illustrative purposes only and is based on estimates and assumptions that management believed to be reasonable at the time such information was prepared; it is inherently subject to significant business, economic, and competitive uncertainties, risks, and contingencies, many of which are beyond the Company's control. Actual results may differ materially from such information, and the inclusion of such information in this presentation should not be regarded as a representation by any person that the results reflected therein will be achieved. Such information should not be relied upon as necessarily indicative of future results.

Non-GAAP Financial Measures

This presentation refers to Adjusted EBITDA, a non-GAAP financial measure not calculated or presented in accordance with GAAP, which is supplemental information and in addition to the financial measures presented herein that are calculated and presented in accordance with GAAP.

EBITDA is defined as net income (loss) before interest expense, net; income tax expense (benefit); and depreciation and amortization. Adjusted EBITDA is defined as EBITDA further adjusted to exclude stock-based compensation expense and gains and losses arising from the issuance, remeasurement, extinguishment or disposition of financial instruments, investments and other assets. Management uses Adjusted EBITDA to evaluate operating performance, establish budgets and forecasts and make operational decisions, and the company believes the measure is useful to investors because it excludes items that management does not consider indicative of the underlying operating performance of the business and because it facilitates comparison of results across periods. Adjusted EBITDA has limitations as an analytical tool and should not be considered in isolation or as a substitute for net income (loss) or any other measure prepared in accordance with GAAP. Among other limitations, Adjusted EBITDA excludes depreciation and amortization of assets that are central to the company's ability to generate revenue, and excludes interest expense and amortization associated with the company's finance lease obligations, which are quantified in the company's earnings releases and filings with the SEC. Adjusted EBITDA as defined by the company may not be comparable to similarly titled measures reported by other companies. A reconciliation of net income (loss) to Adjusted EBITDA is included in this presentation and in the company's Q2 2026 earnings release.

Market and Industry Data

This presentation includes information concerning economic conditions, the Company's industry, the Company's markets, and the Company's competitive position, based on a variety of sources, including information from independent industry analysts and publications, as well as QumulusAI's own estimates and research. The Company's estimates are derived from publicly available information released by third-party sources, as well as from its internal research, and are based on such data and the Company's knowledge of its industry, which the Company believes to be reasonable. Any independent industry publications used in this presentation were not prepared on the Company's behalf. This information involves many assumptions and limitations, and you are cautioned not to give undue weight to these estimates. Certain monetary amounts, percentages, and other figures included in this presentation have been subject to rounding adjustments; accordingly, figures may not be exact arithmetic aggregations of the figures that precede them.

Regulation FD

To the extent this presentation is furnished to the SEC under cover of a Current Report on Form 8-K pursuant to Item 7.01 (Regulation FD Disclosure), the information contained herein shall not be deemed "filed" for purposes of Section 18 of the Exchange Act or otherwise subject to the liabilities of that section, and shall not be deemed incorporated by reference in any filing under the Securities Act or the Exchange Act, except as expressly set forth by specific reference in such filing.

// SECTION 01

Why Now, Why Us

// THE OPPORTUNITY

Demand for AI infrastructure providers is compounding.

It all runs on AI infrastructure
that's the opportunity.

Demand is proven; the binding constraint is capacity.

~7×

Google's monthly tokens leaped from 480T to 3.2 quadrillion year over year (May '25–May '26); API throughput rose another 38% QoQ in Q2.

Usage is proven

~3×

Projected rise in US datacenter power demand by 2030, to 180+ GW.

Demand for power

\$720B+

Big-4 hyperscaler 2026 capex guidance — raised for the third straight quarter, up from ~\$410B in 2025.

Capital is ready

\$7T

Global datacenter build-out through 2030.

The infrastructure opportunity

Sources: Google I/O (May 2026) and Alphabet Q2 2026 earnings (Jul 2026) — tokens; Amazon, Alphabet, Meta, Microsoft Q2 2026 earnings (Jul 2026) — FY2026 capex guidance; McKinsey Global Institute (2025) — data-center build-out; S&P Global — surging US datacenter power demand tests sustainability targets (May 2026).

// THE OPPORTUNITY

Drivers of Demand

Sources: Stanford University (Apr 2026) — AI Index Report; Harvard Business Review (June 2026) — How People are Really Using AI in 2026.

Better models.
Better results.
Real use cases.

~100%

of the human baseline on SWE-bench coding tasks, up from 60% one year earlier.

Capability is accelerating.

53%

of the global population adopted Gen AI in three years, faster than the PC or the internet.

AI adoption is spreading.


AI USE CASES



Personal and Professional Support




Customer Service/
Call Center



Software Development



Technical Assistance
and Troubleshooting



Content Creation
and Editing



Research, Analysis,
and Decision-making

// THE DEMAND SIGNAL

Demand is real and it's repricing compute

Sources: Exponential View, State of the AI Economy (Jun 2026) — GenAI revenue; Epoch AI, An update on AI's most important number (Aug 2026) — OpenAI/Anthropic run rate; CBRE, North America Data Center Trends H1 2026; CoreWeave Q2 2026 earnings call (Aug 11, 2026) — A100 recontracting; SemiAnalysis H100 1-Year Rental Price Index (Apr 2026). Chart: SemiAnalysis AI Cloud TCO Model & GPU Pricing Product. Values approximate.

\$175B+

annualized GenAI revenue (ex-China); OpenAI + Anthropic alone doubled in ~4 months

80%


datacenter capacity under construction already pre-leased — **sold out before it's built**

2020 → 2029

CoreWeave recently contracted 2020-vintage A100s out to 2029 at an "attractive price"

3Y CONTRACT :: \$ / HR / GPU

GPU Rental Pricing Industry Trends



// THE BOTTLENECK

Data centers can't come online fast enough

Sources: Lawrence Berkeley National Lab, Queued Up 2026 Edition (Jun 2026); Sightline Climate via Bloomberg (May 2026) — 2026 US datacenter pipeline; Gartner (2025).

GRID INTERCONNECTION

5+ years


Median wait from interconnection request to commercial operation — and still lengthening.



2026 PIPELINE

30-50%


of US datacenter capacity planned for completion in 2026 is expected to slip or be canceled.



STUCK IN QUEUE

2,060 GW


Generation & storage awaiting grid connection — 1.5× the entire installed US fleet.



// OUR FACTS MODEL


The market needs faster, more flexible access to compute

We house the latest models of NVIDIA GPUs in colocation facilities and <50MW capacity data centers to bring GPU capacity online in months, on a quarterly deployment cadence that can accelerate time to value for our clients.




Flexibility

Flexibility means control — bare metal or virtualized, reserved or burst — all within a unified framework. Infrastructure that adapts as fast as you do.




Access

Access means availability without compromise. We deliver enterprise-grade capacity when it's needed – through reserved clusters and on-demand pools of next-gen GPUs.




Cost

Cost means pricing that supports growth. AI infrastructure should enable innovation—not limit it. Our approach delivers predictable, competitive economics so customers can secure the compute they need while aligning infrastructure investments with the realities of their business.



Trust

Trust means end-to-end control across the infrastructure we operate — with SLAs, isolation, and privacy built in. We don't just deliver capacity; we deliver reliability, because AI progress depends on confidence.



Speed

Speed means adaptability at scale. From provisioning to expansion, our modular infrastructure keeps development cycles moving – so innovation rarely waits for capacity.

// SECTION 02

The Proof

// READING THE MEGAWATT

What’s inside one megawatt of IT load

Sources: NVIDIA DGX B300 specifications (14.5 kW, 8 GPUs per system); SemiAnalysis datacenter and AI cloud research; QumulusAI management estimates for equipment cost and revenue per MW ranges, representative of market averages as of the date on the cover. Assumes ~10% of IT load for networking and storage and ~10 kW 8-GPU H100 systems; configurations, pricing and utilization vary.

BLACKWELL - CLASS DEPLOYMENT · NVIDIA SPECIFICATIONS AND INDUSTRY AVERAGES · ILLUSTRATIVE



THE MATH

Revenue per MW $\approx$ GPUs per MW $\times$ price per GPU-hour $\times$ 8,760 hours

~ 552 GPUs $\times$ $\sim \$4.30$ $\times$ 8,760 hours $\approx$ \$20.8M a year. Reserved capacity is billed every hour, whether or not the customer is running a job.

WHY WE COUNT MEGAWATTS, NOT GPUS


H100 generation		~ 758 GPUs per MW
Blackwell generation		~ 552 GPUs per MW

Newer GPUs: fewer per megawatt, more revenue per megawatt. A GPU count alone doesn’t tell you capacity or revenue.

// POWER FOOTPRINT

What’s live, what’s signed, and what we hold

POWER BY STATUS



8 MW

Fully sold, coming online

Five colocation sites · Marietta, Kansas City, Denver, Philadelphia (2)

3.75 MW

Contracted for Q4

Metro Atlanta, 7-year colocation
ROFO on 7 MW more if the site expands

39 MW

Powered land, optionality

Oklahoma leased up to 2044


// CAPACITY INTO REVENUE

What today's capacity generates, and its potential

WHAT EACH BLOCK OF OUR CAPACITY IS WORTH PER YEAR

3.75 MW
Atlanta · Q4 2026
contracted colocation

8 MW
Five sites · Live 2026
contracted colocation



~\$75M / yr

Potential revenue, at ~\$20M per MW, once sold and fully active — within the ~\$19–21M per MW range we see in the market

~\$132M / yr

Fully sold, at ~\$16.5M per MW. With all deployments expected to come online by year-end 2026

~\$207M / yr


With both blocks sold and deployed

11.75 MW across six colocation sites. ~\$17.6M blended revenue per MW across the blocks. Illustrative annual billing at full deployment only. Not a guarantee of future prices. Represents management’s best estimate as of the date on cover.

// REVENUE BACKLOG

Q2 RPO of \$173.1M; ~\$276M pro forma backlog with post-quarter signings

EXPECTED REVENUE RECOGNITION BY YEAR – REPORTED Q2 RPO VS. PRO FORMA BACKLOG (ILLUSTRATIVE)



+59%

Since Q2 2026 quarter-end

Left: RPO as reported in the Q2 10-Q — \$34.4M expected to be recognized in 2026, \$138.7M in 2027–2029.

Right: the same RPO plus ~\$103M of contracts signed between quarter-end and the Aug 25 filing. We assume ~5% (~\$5M) of the new signings is recognized in H2 2026; the remaining ~\$98M joins the 2027–2029 layer.

// TRACTION & PROOF POINTS

Selected news, June–August 2026

AUG 20, 2026	Metro Atlanta site anchors the next 3.75 MW Seven-year colocation, RFS Q4 2026; right of first offer on 7 MW more — a path to 10.75 MW.
AUG 11, 2026	DRW becomes a customer for a dedicated B300 cluster Annually renewable up to four years with the global trading firm — a second financial-markets customer.
AUG 7, 2026	Agentic hedge fund runs on QumulusAI Blackwell capacity Reserve capacity earns market-rate compute plus a share of trading profits, with no exposure to losses.
JUL 28, 2026	\$71.9M three-year agreement with an AI inference platform Among the largest commitments to date; B300 and B200 capacity ready for customer use in Q3 2026.
JUL 23, 2026	\$32M two-year Blackwell B300 agreement for generative media Dedicated clusters for an inference platform serving image and video workloads; capacity online in fall 2026.
JUL 22, 2026	\$18M two-year take-or-pay agreement with a GPU marketplace Makes QumulusAI a core Blackwell B300 supplier to a cloud marketplace serving AI teams in 100+ regions.
JUN 11, 2026	\$124M+ in three-year inference agreements Hyperbolic and a leading inference platform — 1,280 NVIDIA Blackwell GPUs, ~\$21.9M in upfront commitments.

// DISTRIBUTED NETWORK

Distributed for speed now, positioned for inference next

6

Metro-adjacent colocation sites

Atlanta, GA
Marietta, GA
Kansas City, MO
Denver, CO
Philadelphia, PA – x2

2

Powered-land sites


Oklahoma and Texas · long-term land leases with
~39 MW of grid power available today — optionality
for future HPC capacity development.

1

Pilot-to-production IXP — Wichita State, Q4 2026

Pilot with CNIXP. On success, scale to ~25 Phase-1 sites, with a
path to 125; subsequent sites selected as pilots prove out.

 Contracted colocation  Powered land  AI IXP pod · CNIXP partnership



PARTNERSHIP · CNIXP

GPU pods at the edge

Compute at carrier-neutral Internet
Exchange Points where networks
already meet — fewer hops, lower
latency. Wichita State is the pilot-to-
production site; further sites follow in
phases after successful pilots, as the
partner delivers land, power and shell.

// SITE BUILD MODEL

A network of sites: how we build it, and why it compounds

HOW WE BUILD

Find

Pockets of 2–50 MW of existing power.


Deploy

Racks in, GPUs live — in months, not years.

Monetize

The goal: a shorter path to revenue.

WHAT WE’VE BUILT



COLOCATION

Metro-adjacent colocation


Existing capacity in and around major metros.



POWERED LAND

Power we hold ourselves

Long-term land leases, grid power in hand.



PARTNER EDGE

GPU pods at the edge

Small pods where networks already meet — Wichita State pilot first, scaling on success.

“Today, distribution is a means to an end.
Over time, we think it becomes the end itself.”

MICHAEL MANISCALCO, CHAIRMAN & CEO
Q2 2026 EARNINGS CALL

// TRACTION & PROOF POINTS

Who We Work With

BLUE CHIP PARTNERS

































BREAKOUT CLIENTS
























// SECTION 03

Why We Win


// WHY WE WIN

Customers don't open with price. They open with "When can I be running?"

TIME TO A TOKEN-GENERATING CLUSTER

HYPERSCALE BUILD

3-4 years
to revenue generation



Speed

decides most deals today

We believe supply is constrained from hyperscalers to startups — and that when supply can't keep pace, speed stops being a feature and becomes the decision. Customers buy five things: flexibility, access, cost, trust and speed. All five matter; in our view, one of them decides today.

Sources: JLL, Global Data Center Outlook (2026) — grid connection wait times; industry build-timeline benchmarks for hyperscale / AI-grade greenfield projects (permitting 6-18 mo.; design & development 6-12 mo.; construction 24-36 mo.). Illustrative; actual timelines vary by market and project.

// COMPETITIVE DIFFERENTIATION

Smaller, distributed sites reach power faster than large campuses.

Pockets of 2–50 MW of existing or near-ready power, mostly in colocation, brought online in months.

Why smaller and distributed sites win on time

1. More sites to choose from, and in many jurisdictions, a lower bar for permits and approvals.
2. Faster builds and lower delivery risk than a multi-year campus.
3. The right size for most customers outside frontier-model training.
4. Shorter contracts by design, so part of the book reprices each year.



// NVIDIA CLOUD PARTNER

Validated by NVIDIA. Built for production.

QumulusAI is an official NVIDIA Cloud Partner, recognized as delivering infrastructure purpose-built for modern AI workloads at production scale. Our platform is built and validated against NVIDIA reference architectures.

01

Faster time to production

Reference-architecture builds, not bespoke bring-up.

02

Consistent performance at scale

Validated fabric, cooling and power — every cluster.

03

Confidence in the investment

Proven accelerated computing, end to end.



"Becoming an NVIDIA Cloud Partner **affirms what our customers already experience** — enterprise-grade NVIDIA compute, delivered at hyperspeed. It's about giving customers what they need today, while building toward a larger inference future we believe is coming."

— Michael Maniscalco, CEO, July 17, 2026

NVIDIA PARTNER NETWORK [NVIDIA.COM](https://www.nvidia.com/partner-network) →

// FROM THE SOURCE

Customer testimonials

"QumulusAI's workload-optimized infrastructure gives us the performance, efficiency, and scalability we need as we continue expanding reliable GPU compute for customers building AI at scale."

— Jasper Zhang, CEO of Hyperbolic

"This capacity allows us to run more complex research at greater scale, iterate faster and shorten the distance between a hypothesis and what the data tells us. That ability to keep learning is fundamental to finding new edge."

— Rich Norman, Chief Information Officer, DRW

"This deployment gives AI inference platforms on Shadeform access to dedicated, enterprise-grade AI infrastructure as they scale their businesses."

— Ed Goode, chief executive officer, Shadeform



// SECTION 04

Team & Governance

// TEAM & GOVERNANCE

At the helm: an operator who's scaled GPU cloud

Michael Maniscalco

Chief Executive Officer :: QumulusAI

Michael has built and launched a GPU-as-a-Service business at scale before — the exact business QumulusAI runs today.



01 — FOUNDATION

Georgia Tech

Computer science degree.

02 — VENTURE

Stanley X

Entrepreneur-in-Residence.

03 — GPUAAS AT SCALE

Applied Digital (CTO)

Launched the GPUaaS solution.

04 — TODAY

QumulusAI (CEO)

Leading since September 2025.

// TEAM & GOVERNANCE

A team that has built this before



Mike Maniscalco

CHIEF EXECUTIVE
OFFICER + DIRECTOR



Scott Krosnowski

CHIEF FINANCIAL
OFFICER



Ankur Chatterjee

CHIEF INTEGRATION
OFFICER



Ryan DiRocco

CHIEF TECHNOLOGY
OFFICER



Stephen Hunton

CHIEF MARKETING
OFFICER



Patrick Gahan

SENIOR VICE PRESIDENT
CAPITAL MARKETS + DIRECTOR



Jim Marino

SENIOR VICE PRESIDENT
SALES



Steve Gertz

CHIEF GROWTH OFFICER
FORMER CHAIRMAN



Zachary Graeve

HEAD OF INVESTOR
RELATIONS



Stacy Kenworthy

LEAD INDEPENDENT
DIRECTOR



Michael Mulica

DIRECTOR



David Rench

DIRECTOR
COMPENSATION CHAIR



Barry Schwartz

DIRECTOR
NOMINATING CHAIR



Homaira Akbari

DIRECTOR
AUDIT CHAIR

// SECTION 05

Financial Overview

Q2 2026 financial summary

¹ Adjusted EBITDA is a non-GAAP measure; see the reconciliation of net income (loss) to Adjusted EBITDA in this presentation and in the Q2 2026 earnings release at investors.qumulusai.com.

\$ in actual dollars · Three months ended June 30, 2026

	Q2 2025	Q2 2026	YoY Δ
Revenue	3,085,439	6,712,805	+3,627,366
Cost of revenue	1,386,374	2,242,402	+856,028
Gross profit	1,699,065	4,470,403	+2,771,338
Operating loss	(2,181,533)	(7,670,591)	(5,489,058)
Net income (loss)	12,118,574	(22,776,206)	(34,894,780)
Adjusted EBITDA (non-GAAP) ¹	(266,136)	(782,403)	(516,267)

\$173.1M

Remaining performance obligations

as of June 30, 2026; contracted revenue not yet recognized

66.6%

Gross margin

up from 55.1% a year ago; GPU activations grew revenue faster than colocation costs

+96%

Revenue vs. Q1 2026

sequential growth — +118% year over year

\$(0.8)M

Adjusted EBITDA ¹

near breakeven beneath the non-cash convertible-note items

Statement of Operations

¹ Q2 2026 includes a \$19.2M non-cash loss on issuance of the ATW convertible note, partly offset by \$6.2M of fair-value gains; H1 2026 includes \$73.9M of such losses. Q2 2025 includes a \$14.5M one-time TCM remeasurement gain.

GAAP · \$ in actual dollars · Three and six months ended June 30, 2026

+118%

Q2 2026 revenue growth (YoY)

+104%

H1 2026 revenue growth (YoY)

\$5.6M

Q2 2026 compute revenue, 84% of total (up from 43%)

	Q2 2025	Q2 2026	H1 2025	H1 2026
Revenue from cryptocurrency mining	148,038	410,081	295,839	779,060
Revenue from mining hosting services	1,625,701	693,778	3,349,152	1,642,205
Revenue from compute power	1,311,700	5,608,946	1,311,700	7,711,458
Total revenue	3,085,439	6,712,805	4,956,691	10,132,723
Cost of revenue	1,386,374	2,242,402	2,601,468	4,378,199
General and administrative	2,454,605	4,342,620	3,263,676	8,250,275
Sales and marketing	366,093	916,220	664,293	1,152,572
Depreciation and amortization	1,059,900	6,882,154	1,537,948	9,548,999
Total costs and expenses	5,266,972	14,383,396	8,067,385	23,330,045
Operating loss	(2,181,533)	(7,670,591)	(3,110,694)	(13,197,322)
Total other income (expense), net ¹	14,585,227	(15,168,769)	13,691,911	(59,380,659)
Income (loss) before income taxes	12,403,694	(22,839,360)	10,581,217	(72,577,981)
Income tax expense (benefit)	285,120	(63,154)	285,120	(185,064)
Net income (loss)	12,118,574	(22,776,206)	10,296,097	(72,392,917)

Balance sheet & liquidity snapshot

Total long-term liabilities at June 30, 2026 include convertible notes at fair value (\$55.5M) and the additional convertible notes option liability (\$38.7M), both non-cash fair-value items; finance lease liabilities of \$45.8M; and USD.AI protocol loans of \$18.9M.

\$39.9M

Cash and restricted cash at June 30

\$215.0M

Total assets

\$30.5M

Deferred revenue — customer prepayments

Balance Sheet

	Dec '24	Dec '25	Jun '26
Total current assets	4,785,983	12,905,233	53,204,470
Property and equipment, net	2,737,019	12,502,886	44,006,002
Goodwill & intangible assets, net	—	38,685,340	38,644,691
Other non-current assets	12,822,892	27,641,034	79,147,053
Total assets	20,345,894	91,734,493	215,002,216
Total current liabilities	8,582,104	12,140,082	66,507,183
Total long-term liabilities	10,321,889	14,302,280	149,276,297
Total liabilities	18,903,993	26,442,362	215,783,480
Total stockholders' equity (deficit)	(590,660)	65,292,131	(781,264)

Cash Flow / Liquidity

	FY24	FY25	H1'25	H1'26
Operating activities	(2,021,408)	(5,677,062)	(795,882)	22,305,214
Investing activities	(1,871,322)	(8,594,008)	3,702,653	(36,330,819)
Financing activities	7,231,852	22,013,097	(2,896,954)	42,205,404
Net change in cash	3,339,122	7,742,027	9,817	28,179,799
Ending cash & restricted	3,970,466	11,712,493	3,980,283	39,892,292

// OPERATING LOSS → NET LOSS BRIDGE

H1 2026 net loss is almost entirely non-cash


Six months ended June 30, 2026, \$ in millions.

¹ Non-cash. Fair-value gains are on the convertible notes and additional-notes option liabilities.

² Interest expense, net; warrant liability fair-value change; income tax benefit; and all other items between operating loss and net loss.

Operating loss was \$13.2M for the H1 2026.

The reported \$72.4M net loss is driven by **\$73.9M of non-cash losses** on issuance of the two convertible note tranches, partly offset by a \$12.6M gain on the sale of equity-method investments — not by operations.



Adjusted EBITDA reconciliation

Adjusted EBITDA is a non-GAAP measure with limitations as an analytical tool and is not a substitute for net income (loss); see the Q2 2026 earnings release for the full definition, reconciliation and limitations, including finance-lease payments it excludes.

\$ in actual dollars · Three months ended June 30, 2026

Q2 2026 ADJUSTED EBITDA

\$

(0.8)M

Q2 2025: \$(0.3)M

Revenue

\$

6.7M

vs. \$3.1M in Q2 2025

Operating loss

\$

(7.7)M

vs. \$(2.2)M in Q2 2025

Adjusted EBITDA is EBITDA further adjusted to exclude stock-based compensation and gains and losses arising from the issuance, remeasurement, extinguishment or disposition of financial instruments, investments and other assets. Management uses it to evaluate operating performance; it excludes items not considered indicative of the underlying business.

RECONCILIATION	Q2 2025	Q2 2026
Net income (loss)	12,118,574	(22,776,206)
+ Income tax (benefit)	285,120	(63,154)
+ Interest expense, net	601,260	2,045,748
+ Depreciation and amortization	1,059,900	6,882,154
EBITDA	14,064,854	(13,911,458)
+ Stock-based compensation	233,208	119,939
+ Loss on issuance of convertible notes	—	19,241,000
+ Change in fair value — convertible note	—	(2,380,000)
+ Change in fair value — additional notes option	—	(3,850,850)
+ Gain on remeasurement of TCM investment	(14,549,536)	—
+ Other, net	(14,662)	(1,034)
Adjusted EBITDA (non-GAAP)	(266,136)	(782,403)

// WE ARE THE INFRASTRUCTURE THAT RUNS AI.

This isn't hyperscale.
It's hyperspeed.