



SOAR

SOVEREIGN AI RESOURCES · NYSE AMERICAN: SOAR

Investor presentation

Secure intelligence. Sustainable infrastructure.

October 2026

Disclaimer and forward-looking statements

This presentation has been prepared by Volato Group, Inc. (NYSE American: SOAR), which is changing its corporate name to SOAR AI, Inc. (together with its subsidiaries, the “Company”), and is being provided on a confidential basis solely to prospective investors that have agreed to keep its contents confidential. This presentation does not constitute an offer to sell or a solicitation of an offer to buy any securities. Any offering will be made only to qualified investors by means of definitive documentation.

This presentation contains forward-looking statements within the meaning of the Private Securities Litigation Reform Act of 1995, including statements about the development, construction, power, cost and timing of the Company’s data center capacity; the deployment and acceptance of AMD Instinct MI355X and MI455X GPU capacity; contracted and expected revenue and customer prepayments; customer pipeline; financing plans, including the proposed private placement, senior debt and an equity facility; expansion to additional sites; and the Company’s name change and listing plans. Words such as “anticipate,” “expect,” “plan,” “target,” “intend,” “believe,” “estimate,” “may,” “will” and similar expressions identify forward-looking statements.

Actual results may differ materially because of risks and uncertainties, including: construction, equipment delivery and commissioning delays or cost overruns; power and cooling availability; GPU supply and allocation; customer delivery, acceptance, service-level and performance requirements; concentration of contracted revenue in a single customer; the availability and terms of equity, convertible and debt financing; dilution; integration of the recently completed merger; stockholder and NYSE American approvals and continued listing; and the other risks described in the Company’s filings with the U.S. Securities and Exchange Commission, including its most recent Annual Report on Form 10-K and Quarterly Reports on Form 10-Q. Forward-looking statements speak only as of the date of this presentation, and the Company undertakes no obligation to update them except as required by law.

Contract values shown are aggregate contractual amounts over the full term of the applicable orders and are not projections of recognized revenue. Capacity, job and investment figures are management estimates and projections. AMD, AMD Instinct, ROCm and Helios are trademarks of Advanced Micro Devices, Inc.; other names are used for identification only.

What has changed

Recent developments

The merger has closed, customer contracts are signed, and construction of the first campus is underway.

Merger closed SEP 11 Alignment Engine merger completed September 11, 2026. The AI infrastructure business now operates inside a NYSE American-listed company.	\$1.17B contracted SEP 22 Master services agreement and two four-year take-or-pay orders signed with a large AI customer: 4,000 MI355X and 4,000 MI455X GPUs.
Capacity backstop in place MAY 2026 Backstop of approximately \$485M covering unused Phase I capacity, signed in May 2026, reduces exposure to placing that capacity.	Construction underway IN PROGRESS Ohio Phase I build mobilized in September. Initial customer service is targeted for December 31, 2026.
SOAR AI OCT 8 Rebrand announced October 8. Legal name change to SOAR AI, Inc. expected October 19, with trading under the new name expected October 20.	Up to \$1B equity facility OCT 12 Announced October 12. The Company decides whether, when and how much to sell, and draws as growth requires, alongside prepayments and project-level financing.

\$1.17B

Contracted orders

Two 4-year take-or-pay orders

8,000

GPUs under contract

MI355X and MI455X

154 MW

Ohio power available

Roadmap to 480 MW

~\$133M

Customer prepayments

Under signed orders

Contracted demand before first service.

Take-or-pay orders anchor both deployment phases and support project financing.

Powered industrial campus.

Former industrial site with existing power, permits, water supply and wastewater treatment shortens the path to capacity and local approval.

AMD Instinct platform.

MI355X today, MI455X and Helios rack-scale systems next, on the open ROCm software stack.

Phased, financeable build.

Each phase is funded by its own mix of equity, senior debt and customer prepayments.

Repeatable model.

Repurposed industrial and greenfield sites with behind-the-meter power are being evaluated in the U.S. and internationally.

Public-company platform.

NYSE American listing and an up to \$1B equity facility provide access to capital as capacity scales.

SECTION 01

Business overview



We develop powered AI data centers and provide GPU capacity to customers that need control over their data and models.

AI infrastructure

- Develop and operate AI data center capacity, starting with the Ohio campus
- Dedicated and shared GPU capacity, plus on-prem and edge systems
- Signed orders are dedicated, single-tenant capacity under long-term contracts
- Proprietary NeuralSync data movement technology

Community-integrated campuses

- Repurposed industrial sites, or greenfield sites with behind-the-meter power
- No draw on the municipal water supply at the Ohio campus
- Local hiring and manufacturing jobs, with commitments the host town can verify


Customer control

- Dedicated hardware in a specified location, with audit rights
- Customers control their data, models, logs and the physical location of their hardware
- Open AMD ROCm stack lets customers choose their models and tools

GPU-as-a-Service business model

How we earn revenue

Customers reserve dedicated compute capacity for a fixed term and pay whether or not they use it.



Why it works for investors. Fees begin when capacity is delivered and accepted, and are subject to service-level and performance requirements. Matching each deployment phase to signed orders before major procurement keeps capital tied to contracted revenue.

Contracted revenue today comes from dedicated capacity. Additional offerings broaden the customer base over time.

TODAY

Dedicated GPU capacity

- Single-tenant AMD Instinct clusters reserved for one customer
- Long-term take-or-pay contracts with prepayments
- MI355X service from December 31, 2026; MI455X from June 30, 2027

NEXT

Shared capacity

- Training, inference and enterprise AI on shared infrastructure
- Serves customers below the size of a dedicated cluster

NEXT

On-prem and edge systems

- NeuralSync-based systems installed in the customer's own facility
- For data that must stay on site

NEXT

Next-generation platforms

- Helios rack-scale systems on the AMD roadmap
- Infrastructure and platform services alongside compute

AMD Instinct platform and technology

Technology

We build on AMD's GPU roadmap and open software stack, and add our own data-movement technology.

AMD GPU roadmap

- Phase I: AMD Instinct MI355X
- Phase II: AMD Instinct MI455X, with higher memory capacity and bandwidth per GPU
- Future phases planned around the Helios rack-scale platform

Open software

- AMD ROCm open software ecosystem supports leading AI frameworks and models
- Customers choose their own models, frameworks and tools
- Portability lowers the cost of switching hardware generations

NeuralSync DPU

- Proprietary FPGA-based network accelerator
- Designed to reduce network bottlenecks so GPUs spend less time waiting for data
- Patented technology; boards can be built for on-premises systems

Supply chain. Supermicro supplies the GPU server systems, and DDN supplies high-performance storage. We integrate the full stack and operate it for the customer.

Where AI infrastructure demand is going

Market

Demand is shifting toward dedicated capacity, and communities are scrutinizing how new data centers are built.

56% → 41%

Public-cloud share of AI inference

One-year change (Broadcom)

\$810.6B

AI data center market by 2033

From \$147.3B in 2025 (Grand View)

45 projects

Blocked or delayed in Q2 2026

by local opposition, \$68B

Customers we serve

- **AI platforms and model builders:** large, long-term blocks of dedicated GPUs
- **Regulated enterprises:** banks, insurers, health systems needing control and audit trails
- **Public sector:** state, local and federal buyers needing data residency
- **Host communities:** reserved private capacity under the community agreement

Why we can win

- Dedicated capacity in a specified location, with audit rights
- Sites with their own power or existing infrastructure draw fewer local objections
- Open AMD stack and commercial inference platforms make private AI practical
- Contracted capacity before construction is complete

Sources: Broadcom survey of 1,800 IT leaders (Feb.–Mar. 2026), vendor-sponsored; Grand View Research, global AI data center market, 23.9% CAGR 2025–2033, not independently verified by the Company; Data Center Watch via Tom's Hardware.

SECTION 02

Contracted demand



Contracted demand | \$1.17 billion of take-or-pay orders

Customer contracts

A master services agreement and two four-year orders with a large AI customer cover both deployment phases.

Order #1 AMD Instinct MI355X	
GPUs	4,000
Contractual start date	December 31, 2026
Term	48 months
Contract value	\$399.5M
Customer prepayment	\$39.9M


Order #2 AMD Instinct MI455X	
GPUs	4,000
Contractual start date	June 30, 2027
Term	48 months
Contract value	\$773.7M
Customer prepayment	\$92.8M

Contracted run-rate. Approximately \$24.4M per month, or about \$293M a year, once both orders are in service. Fees are take-or-pay and are subject to delivery, acceptance, service-level and performance requirements.

Contracted revenue profile

Signed orders only

Annual contract revenue under the two signed orders, assuming capacity is accepted on the contractual start dates.



\$1,173M
Total contract value
48-month terms, 2027 to mid-2031

\$293M
Full-year contracted revenue
2028 to 2030, both orders in service

Upside Not Shown

Growth beyond signed orders

- Unused Phase I capacity, covered by the capacity backstop
- Campus expansion and additional sites
- Other customers now in discussion

SECTION 03

AI infrastructure



Building 285 at a former paper mill is being converted into a liquid-cooled AI data hall.

9,216

AMD MI355X and MI455X GPUs

25 MW

Phase I design block

49 MW

Lease power by Feb. 2027

~\$730M

Phase I program incl. compute

240

Projected jobs

Phase I build

- 9,216 GPUs on 1,152 Supermicro servers across 197 racks
- Major equipment self-procured; CK Construction is construction manager at risk
- Liquid cooling on existing mill water infrastructure; chiller plant adds full capability in 2027

Why an industrial site

- Reuses existing buildings and industrial infrastructure on a large mill property
- Cooling water from existing industrial wells, separate from the municipal supply
- Council lifted its moratorium on September 14 after a public review session


Room to grow

- Ohio site: 154 MW available today, with a roadmap to 480 MW
- Lease includes expansion rights to additional buildings
- Projected \$740M of equipment added to the local tax base

Path to first service


Phase I and II milestones

Construction is underway against a schedule that targets customer-ready capacity on December 31, 2026.



Schedule management. Long-lead items (power cable, switchboards, servers and chillers) are tracked weekly. Each order carries a contractual start date, and capacity must pass the customer’s acceptance tests before fees begin.

Prove the model in Ohio, scale the campus, then repeat it at additional powered sites.



Discipline. Each phase is committed only when its customer contracts and financing are in place, so growth does not get ahead of contracted demand or committed capital.

How we choose sites

We weigh each potential campus against these criteria, whether a repurposed industrial site or a greenfield site.



A guide, not a checklist. The same framework applies to opportunities in the United States and internationally. A site does not need to meet every criterion; we may pursue a site where the overall opportunity is attractive.

Where we're going

Roadmap

Deliver the contracted capacity, scale the Ohio campus, then replicate the model at new sites.

2026

Deliver

- Order #1 contractual start: December 31
- Close private placement and equipment debt financing
- Name change to SOAR AI, Inc.



2027

Scale

- Order #2 contractual start: June 30
- Lease power steps up to 49 MW
- Place unused Phase I capacity with additional customers



BEYOND

Replicate

- Grow the Ohio campus along the roadmap to 480 MW
- Add sites in the U.S. and internationally
- Broaden offerings: shared capacity, on-prem and edge

Forward-looking. Each step depends on financing, power, equipment delivery and customer contracts, as described in the forward-looking statements.

SECTION 04

Capital plan and offering



Capital plan | Funding by phase

Sources and uses

We build against long-term contracted demand, funding each phase with customer prepayments, debt and equity.

	Uses	Equity / convertible	Senior debt	Customer prepayments	Total sources
Phase I 9,216 GPUs	~\$730M	\$158M	\$450M	\$186M	\$794M
Phase II 4,000 MI455X	\$596M	\$117M	\$400M	\$94M	\$611M
Total initial deployment	~\$1,326M	\$275M	\$850M	\$279M	\$1,404M

Equity and convertible

- ~\$200M private placement now in market
- Up to \$1B equity facility, drawn at the Company's option as growth requires

Senior and equipment debt

- \$450M Senior Debt A for Phase I
- \$400M Senior Debt B for Phase II
- Secured by contracted take-or-pay revenue

Customer prepayments

- \$132.8M under signed orders
- Balance from additional Phase I customers
- Backstop covers unused Phase I capacity

Private placement summary

Proposed terms

Proceeds fund the equity portion of the Phase I and Phase II deployments.

Issuer	Volato Group, Inc. (NYSE American: SOAR), to be renamed SOAR AI, Inc.
Offering size	Estimated at \$200 million
Securities	Convertible notes and common stock
Pricing and conversion terms	[To be set with the lead investor]
Use of proceeds	~\$115M Phase I (MI355X) and ~\$85M Phase II (MI455X)
Stockholder approval	Issuances above 20% subject to approval at a special shareholder meeting
Registration	Resale registration statement following closing
Placement agent	Lucid Capital Markets

\$200M

Estimated placement

\$1,326M

Phase I + II program

Investor protections

- Proceeds allocated by phase
- Phase II spend gated to committed debt and prepayments
- Contracted revenue before first service

Key risks and how we manage them

Execution

The main concerns we have heard from investors, and what is in place to address them.

Construction and schedule

Owner-procured long-lead equipment, construction manager at risk, weekly critical-path tracking against the Dec. 31 start date

Customer concentration

Signed take-or-pay orders cover 8,000 GPUs; backstop covers unused Phase I capacity; additional customers in discussion

Financing

Phase-by-phase funding with debt and prepayments; Phase II procurement gated to committed capital

Listing and structure

Merger closed; stockholder approvals and NYSE American steps sequenced around a special shareholder meeting

Technology roadmap

MI355X today with a path to MI455X and Helios on the open ROCm stack

SECTION 05

Leadership and milestones



CE

Chris Ensey

Chief Executive Officer

- 25 years building mission-critical, hyperscale and cloud systems
- Former Interim CEO and COO of Riot Blockchain (RIOT) and CTO of BlueVoyant
- B.S. Computer Engineering, Virginia Tech

MH

Mark Heinen

Chief Financial Officer

- 30+ years in finance, 25+ in public companies, 10+ as CFO
- Reverse mergers, deSPACs, M&A and convertible financings
- CPA; MBA and BBA, University of Oklahoma

DS

David Stanfill

Chief Technology Officer

- 12 years in high-performance computing
- Six years leading teams that design and run high-density FPGA and GPU data centers
- B.S. Computer Science, Kent State University

MP

Mike Prachar

Chief Operating Officer

- Chief Operating Officer since February 2022
- Previously COO of Big Green IT, a Tier 1 direct Microsoft cloud service provider
- 20+ years leading operations in telecom and technology services, including Rapid Link, Inc. and WorldPort Communications

Upcoming milestones

Next 9 months

Near-term events that will show progress against the plan.

Oct. 19–20, 2026	Name change to SOAR AI, Inc. effective; trading under the new name	Dec. 2026	Special shareholder meeting, including financing and listing proposals
Mid-Nov. 2026	First quarterly report as a combined company (Form 10-Q)	Dec. 31, 2026	Order #1 contractual start: 4,000 MI355X GPUs
Nov. 30, 2026	Internal power-on at the Ohio campus	Feb. 2027	Chiller plant complete; lease power steps up to 49 MW
Q4 2026	Close private placement and equipment debt financing	Jun. 30, 2027	Order #2 contractual start: 4,000 MI455X GPUs

SOAR AI · NYSE AMERICAN: SOAR

Thank you

Investor contact: investors@soarai.com
8050 Freedom Ave NW, North Canton, Ohio 44720

