August 12, 2026

Nebius Group Letter to shareholders Q2 2026

Q2 2026: Demand inflecting

NEBIUS

Demand, proven in signed contracts

ACV per MW is stepping up[1]

Capacity coming online through 2027
All numbers are approximate

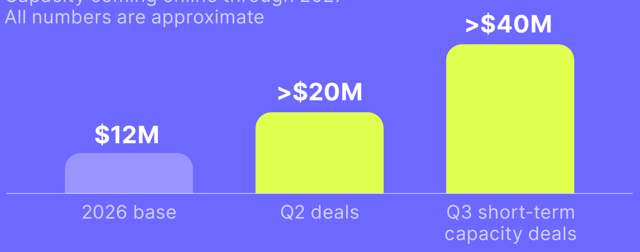

Total contract value grew

4x

quarter over quarter
for deals closed in Q2

Payback period[2][3]

1 year 10 months

for deals closed in Q2

Landmark AI cloud wins in Q2

4 large deals
such as Reflection and Cohere

>$1B average deal size
across the four deals

50-60% of capex self-financed
by prepayments

Capacity footprint

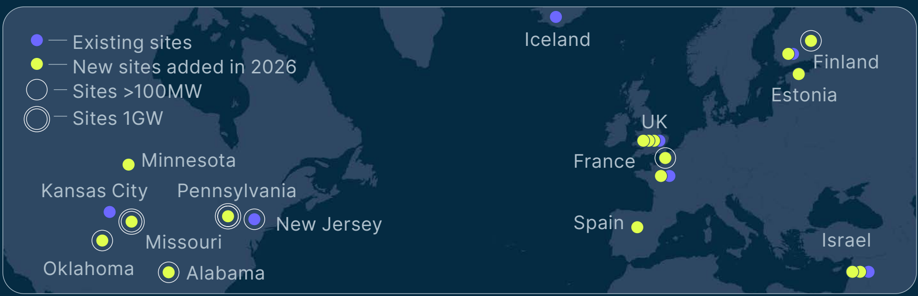

Raising capacity guidance for YE'26

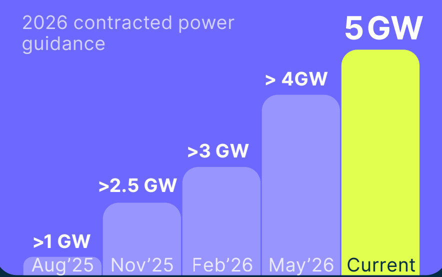

Financial highlights

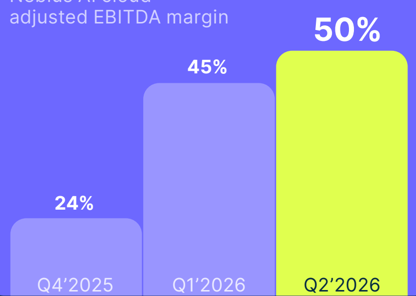

Q2 Nebius group revenue
Up 454% YoY to $582M

2026 FY guidance reiterated

Strong cash position ending Q2
$8B cash
Including $2.3 billion in positive operating cash flow

First asset-backed financing
Repeatable framework
~$775M secured in July
at attractive terms (SOFR + 2.50%)

Building the full-stack AI cloud

Nebius AI cloud "Aether 3.6"

Introduced Nebius Echo — an AI agent to help developers operate and manage workloads

Enhancing the platform to support greater volumes as customer commitments scale

Open ecosystem

New open-weight models in Token Factory:

Kimi K3 | GLM 5.2 | Nemotron Ultra

Token Factory production inference workloads increased >3x in Q2

Eigen AI and Clarifai teams joined Nebius this quarter

Other businesses and investments

Strategic equity investments

ClickHouse
Reported valuation of $15B in January 2026 funding round

Toloka
Data solutions business backed by Bezos Expeditions

Other businesses

AVRIDE
Nearly 3x size of autonomous fleet YTD with 1M+ miles in 2026
Surpassed 60,000 commercial rides on Uber platform

tripleten
Launched AI Systems Engineering programs and expanded its enterprise B2B offering

[1] ACV — Annual contract value.
[2] Methodology is on a revenue recognition basis, excluding any prepayment.
[3] Estimated payback period for new deals signed in Q2, based on forecast costs and contracted future capacity, including capacity not yet built.

Dear shareholders,

This was the quarter the **market validated our strategy**.

Demand for AI capacity continues to grow exponentially, and we are converting that demand into contracted, profitable growth.

Q2 was our strongest commercial quarter yet

- Deals saw an average yield of more than $20 million per megawatt.
- We closed four landmark deals for our AI cloud business, for an average total contract value (TCV) of more than $1 billion each, and a yield of $20-25 million per megawatt.
- TCV of Q2 wins grew nearly 4x quarter-over-quarter, with TCV from new customers growing more than 9x.
- New pricing initiatives in early Q3 — such as our **first-ever auction and short-term capacity deals** — are showing promise. We see a price opportunity in the $40-50 million per MW range and we signed our first one this week.
- Overall deal economics strengthened, with 70% of deals including prepayments, covering 50-60% of the associated capex.
- Together, the expected payback period for the associated capex and related operating costs for Q2 deals is **1 year and 10 months, down from our two-to-three year payback period previously.**
- We could sell our entire 2027 capacity on these terms today. We are deliberately not doing so because we see higher value in retaining some capacity for immediate customer needs.

A contract mix built to capture a rising market

We address customer needs with three types of deals. Each has a different duration, pricing, and role in our business.

Our contract mix can capture this rising market on both ends: longer, higher-priced contracts lock in future revenue, while shorter contracts can capture near-term upside.

- **Shorter-duration contracts**, typically between three and six months, is a new effort, for customers with an acute, time-bound need — priced at a significant premium. One such deal has been signed in Q3.
- **Mid-term contracts** — our core business — now with an average duration of one to three years with the world's most ambitious AI companies.
- **Long-term contracts** with investment-grade customers, which help us finance our buildout faster and more efficiently. The $775 million secured facility we raised in July was on the back of one of these agreements — and with $40 billion in customer commitments, we will do more of this.

Arkady Volozh,
Founder and CEO

"We closed our largest AI Cloud deals on our strongest terms to date, at prices that represent a step-change in the economics of our business."

Capacity build-out accelerating, expecting to deploy >1 GW per year in 2027

- We continue to secure power ahead of our expectations, and today we are raising our year-end contracted power target again to 5 GW.
- We expanded our footprint with additional sites and progressed on construction at our owned AI factories in the US.
- The pace in which we bring capacity to market will accelerate as we plan to deploy more than 1 GW per year — and we plan to do so starting in 2027.
- We are innovating in how we build. This quarter we introduced an asset-light partnership model that addresses the two constraints in this industry: capital and capacity.

Software platform for the full AI lifecycle

- Inference workloads on Token Factory more than tripled in Q2, increasingly running on open models.
- Training and inference are converging. Unlike standalone inference providers, Nebius serves the full workload lifecycle on one platform that will drive better total cost of ownership and utilization of compute.
- The acquisitions of Eigen AI and Clarifai brought industry-leading inference optimization in-house.
- We shipped Nebius AI Cloud 3.6 ("Aether 3.6"), scaling the core platform alongside our growing global footprint.

Financial model proving itself as we scale

- **Q2 group revenue** of $582.3 million, was up 454% year-over-year. **Nebius AI cloud** revenue grew 514% year-over-year to $575 million, ARR reached $3.0 billion.
- We also continued to demonstrate the operating leverage inherent in our business: our AI cloud business delivered an adjusted EBITDA margin of 50% in the quarter.
- **We expect over $9 billion in customer prepayments in 2026.** We are reaffirming our full-year 2026 guidance across all metrics.

Arkady Volozh

[4] Annualized run-rate (ARR) is calculated by taking Nebius AI cloud revenue from the last month of the quarter multiplied by 12.

Group Q2 update

Customers: Deal economics strengthened across every dimension

Q2 was our strongest commercial quarter yet:

- **Annual contract value per MW repriced higher.** Our core AI cloud business won deals averaging above $20 million in revenue per MW, driven by increasing prices for new-generation GPUs and more than 30% higher pricing on older-generation GPUs versus Q1.
- **Deal sizes and new logos set records.** We closed four landmark deals, which averaged more than $1 billion in TCV each.
- TCV of deals closed in Q2 grew nearly 4x quarter-over-quarter, with TCV from new customers growing more than 9x.
- **Roughly 70% of deals closed included customer prepayments** — an all-time high.
- Prepayment has become the market standard for securing capacity in a supply-constrained environment.
- **Deal terms shortened the expected payback period [3] on these deals to 1 year and 10 months**, from our historical two-to-three-year range.

These economics will flow through to revenue as the underlying capacity comes online. The majority of these deals were signed against capacity arriving in late 2026, and will contribute primarily to 2027 revenue.

Our contract mix captures this rising market on both ends: longer, higher-priced contracts lock in future revenue, while shorter contracts can capture near-term upside.

In Q3, **we closed our first short-term capacity deal**. This deal represents dedicated capacity for a customer with an acute, time-bound need for compute. The capacity is expected to go live in Q4. Such deals are typically multi-month engagements, priced at a premium, for customers who need to move fast and know exactly what they need.

We also **conducted our first capacity auction pilot**, seeing early success by securing the highest price we have cleared for NVIDIA Blackwell chips to date.

Landmark customer agreements

The strong deals we closed in the quarter included four landmark deals for our AI cloud.

This success in our core market validates our positioning and our ability to continue capturing more market share.

The average TCV for these deals was greater than $1 billion, with annual contract value of $20-25 million per MW.

- **Reflection** selected Nebius for a multi-year agreement to train and run its open-source models on our platform, accelerating open-source AI development while supporting model safety and the broader developer ecosystem.
- **Cohere** is partnering with Nebius to support the growth of its agentic AI solutions for enterprises, grounded in flexible next-generation NVIDIA architectures, technical collaboration, and our track record as a long-term strategic partner.
- **Another US-based AI neolab** chose our platform for frontier model development to accelerate product innovation in response to growing enterprise demand and to support their scaling revenue.
- **A large US-based quantitative trading firm** represents one of our first large-scale customer wins in financial services. This firm chose Nebius for our consistent speed and reliability at scale, as compute performance is critical to trading execution.

These large deals are only a handful of the wins we saw across the range of industries we serve. Customers require infrastructure and tooling that keeps pace with their speed of development and supports their competitive advantage, and are choosing Nebius as their infrastructure backbone.

Some of these key vertical wins include:

- **AMI**, co-founded by Yann LeCun, is developing world models that learn abstract representations of real-world sensor data and make predictions in representation space. AMI uses Nebius compute and storage infrastructure to support its AI research and model training workloads.
- **Basecamp Research** trained EDEN, a 28-billion-parameter biological foundation model, on 9.7 trillion tokens from over a million newly discovered species, entirely on Nebius, at 20x faster training speed compared to their previous internal environments. EDEN exhibits state-of-the-art performance

[3] Estimated payback period for new deals signed in Q2, based on forecast costs and contracted future capacity, including capacity not yet built.

and designs programmable therapeutic molecules validated across different modalities.

- **Prima Mente** scaled its Alzheimer's and Parkinson's epigenetic foundation model from 1 billion to 68 billion parameters on a dedicated Nebius cluster, achieving 89-97% diagnostic accuracy from blood samples.
- **Higgsfield**, one of Nebius's first AI cloud customers, has expanded its usage on the platform by >20x since their first contract. This summer they premiered a 95-minute feature film made entirely with AI on Nebius infrastructure — an industry first that compressed a two-year process into weeks, at a fraction of the cost of traditional filmmaking.

Capacity: Executing our global capacity expansion

The second quarter was once again defined by strong execution and significant progress toward our capacity expansion targets. We continue to bring capacity online to support near-term growth while securing land and power commitments for 2027 and beyond.

Capacity underpinning near-term growth

In Q2, we continued to deliver on our capacity commitments across both AI cloud customers and strategic long-term contracts.

- We brought additional capacity online across sites in the US and Europe during the second half of the second quarter.
- To date, we've **delivered all capacity tranches to Microsoft under the contract**, including two tranches over the past month.

In May, we announced a long-term partnership with **Bloom Energy** to deploy behind-the-meter fuel cells this year. We are excited about this partnership as it adds fuel-cell technology to the range of options available to us for future deployments.

We have already received **our first NVIDIA Vera Rubin NVL72 systems** and are using these early units to validate compute, networking, and orchestration together as a single system. Testing the integration of this new architecture across our full stack will position us to deploy Vera Rubin as a production-ready offering for our customers once shipments scale. We are already engaging with customers for Vera Rubin access.

- To support agentic AI orchestration, tool-calling and memory workloads, we are **complementing our NVIDIA GPU fleet with more general-purpose CPU compute, including early adoption of NVIDIA Vera CPUs.**

Building the foundation for future growth

We continue to secure power ahead of our expectations. **We now anticipate ending 2026 with 5 GW of contracted power[5], up from the +4 GW we indicated last quarter.**

The pace at which we bring capacity to market will accelerate as we **plan to deploy more than 1 GW per year, starting in 2027.**

We expanded our contracted capacity footprint in Q2, with additional sites in the UK, Estonia and Finland. Construction is also underway at some of our owned AI factories in the US.

We also continued to progress the buildout supporting our second Meta agreement, with capacity on track to come online in early 2027.

New asset-light path to capacity

This quarter we introduced an asset-light partnership model to supplement the capacity we bring online ourselves.

Capacity remains a binding constraint on AI adoption, and bringing capacity online at a rapid pace is central to meeting customer demand.

To help unlock this, we launched a model that allows partners to deploy Nebius's full-stack AI cloud platform in their own AI data centers.

Nebius contributes:

- Our full-stack AI cloud software,
- Our systems architecture and reference designs, and
- Our global go-to-market organization that brings the demand.

Under this model, partners get fully-owned AI infrastructure assets, designed to Nebius standards, and a fast route to serve the AI cloud market; Nebius converts partner-financed capacity into high-margin revenue with minimal capital outlay.

This complements our owned and colocated portfolio while reducing the capital burden of adding capacity for customers.

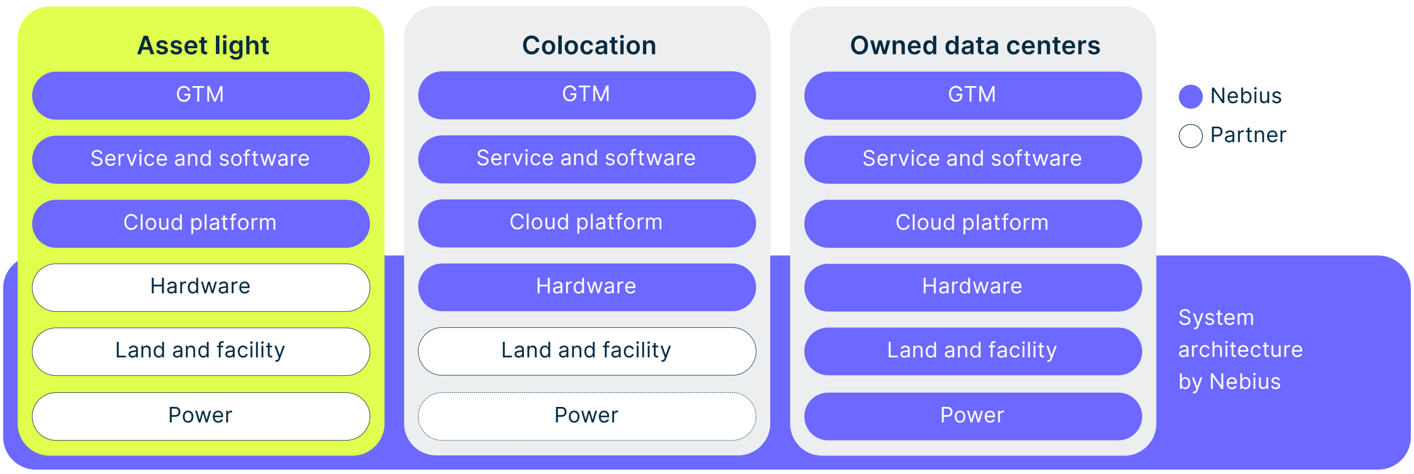

[5] Contracted power, secured by contracted land and power commitments.

Cloud: Advancing an open ecosystem across every layer of the stack

As enterprises put AI into production, inference is becoming one of the fastest-growing layers of AI infrastructure, and increasingly runs on open models for their control, flexibility, and lower cost at scale. Token Factory, our managed inference platform, powers this transition:

Production inference workloads more than tripled in Q2 as customers deployed multiple applications and expanded existing ones. A growing share are agentic (such as search, coding, and customer-facing agents), where a single task drives many model calls, so consumption scales with the complexity of the work automated, rather than with user count.

Training and inference are converging. We see that inference and training use cases are not isolated: Those training models require production inference, while those inferencing open-weight models often evolve toward fine-tuning and post-training as they optimize for quality, latency, and cost.

- Unlike standalone inference providers, Nebius serves the full workload lifecycle on one platform that will drive better total cost of ownership and utilization of compute.
- A growing share of inference demand is generated by training itself — reinforcement learning roll-outs, evaluations, and synthetic data generation.
- A digital health Token Factory customer was able to move their AI mental-wellbeing product to a more capable model and cut worst-case response times by orders of magnitude.

Model performance. Day-one support for the most advanced models like Kimi K3, GLM-5.2, MiniMax 3 and NVIDIA Nemotron 3 Ultra.

- Nebius delivers both accuracy and speed, validated by independent benchmarks.
- In the most recent Endpoint Accuracy Index from Artificial Analysis, **Nebius achieved full accuracy parity on GLM-5.2** while delivering among the fastest output speeds of any provider. The benchmark measures how much of a model's accuracy each inference provider preserves.
- The Eigen AI and Clarifai integrations deepened our model optimization, inference systems, and deployment capabilities.

Agentic infrastructure. Q2 was Tavily's first full quarter within Nebius, pairing Token Factory's inference with the real-time web access that agents need for factual accuracy — reasoning and retrieval on one platform.

- Tavily's developer community has grown to more than 2.5 million developers (from 1 million in February) serving Fortune 500 enterprises and powering production agents.
- In Q2, Tavily launched keyless pay-per-search built for autonomous agent consumption, and achieved ISO 27001 certification for enterprise deployments.
- As model interfaces standardize, differentiation shifts to cost, latency, reliability, quality in production and to the agentic, multi-step workloads defining the next phase of enterprise AI. Token Factory is built for both.

Nebius AI Cloud "Aether 3.6": Strengthening the core platform

In Q2 we shipped Nebius AI Cloud 3.6 ("Aether 3.6"), strengthening the developer experience, governance controls, and storage performance teams need to run AI in production.

As our capacity expands and our customer commitments grow, our cloud platform needs to scale with it. We are enhancing our platform to support greater volume as we close larger customer contracts. We are also deploying our software stack to new regions as we bring on more data center sites globally.

As part of this roadmap, we released the following features as part of Aether 3.6:

- **Developer experience.** The release introduced **Nebius Echo**, our own AI agent for natural-language infrastructure control, built into Nebius AI Cloud and running on open-source models served by Token Factory. Echo lets customers manage their environment in plain language, asking questions grounded in our documentation, checking live resource status and conducting simple operations such as instance creation without commands or additional setup. In the future, Echo will provide automated infrastructure diagnostics and handle complex, multi-step deployments end-to-end.
- **Security and governance.** New controls support sensitive workloads in regulated environments, including a Key Management Service with customer-managed encryption keys, Workload Identity Federation for credential-free authentication, Budgets to help FinOps teams set spending limits, and a Bring Your Own Image option for hardened base images.
- **Storage.** Our storage platform saw several enhancements this quarter. Object storage now delivers 30% higher read speed. Shared Filesystem is 3x faster for small-file operations

Deepening customer collaboration

At our June 9 Inflection Event, we launched our **Customer Advisory Board**, with AMI, Black Forest Labs, Cloudflare, Cognition, Cohere, Core Automation, Higgsfield, Recraft, Revolut, and Rhoda. This forum brings CEOs, CTOs, founders, and industry leaders together to sharpen our customer feedback loop, and keep Nebius at the forefront of AI development.

We also opened the **Nebius Builder Program** in early preview, which provides credits across Nebius AI Cloud, Token Factory, Tavily, and Nebius Academy, with engineering office hours and a builder community. We also launched Nebius Certifications to give AI cloud professionals verifiable credentials.

Expanding executive talent

In July, we welcomed **Lindsey Irvine as our new Chief Marketing Officer.** Lindsey brings deep experience leading global marketing at high-growth technology companies as a multi-time CMO at Square, Benchling and MuleSoft, with prior leadership roles at Salesforce.

Other businesses

Avride

Autonomous vehicle operations achieved important goals in Q2

- The AV-capable fleet has nearly tripled since year-end 2025, exceeding 200 vehicles in May 2026 with more than one million autonomous miles completed this year.
- Avride has completed over 60,000 commercial rides on Uber in Dallas.
- The Dallas operating map has doubled since launch, with further expansion expected.
- The R&D fleet continues to collect mileage and data needed to advance our AI technology stack and accelerate the transition to No-Vehicle-Operator (NVO) operations.

Robodelivery operations launched in Arlington, Virginia, and Miami through UberEats; new Chartwells Higher Education partnership expands campus reach

- Robot deliveries more than tripled year-over-year in Q2 and exceeded 600,000 since inception.
- In Q2, Avride launched robodelivery operations in Arlington, VA and Miami through UberEats.
- In July, Avride announced a Master Services Agreement with Chartwells Higher Education, one of the largest on-campus food service providers in the US. Bowling Green State University is among the first campuses to launch, with more joining this fall.
- Robodelivery operations are live at Indiana University Bloomington.
- Momentum is building with vendor partners following a competitor's exit from campus deliveries.

TripleTen

TripleTen, our edtech business, continues to develop new offerings to meet evolving market opportunities

- TripleTen added AI Systems Engineering as a new program in Q2, continuing to build out a comprehensive suite of AI coursework.
- Tested new products in the US and expanded its enterprise-focused B2B offering, with an emphasis on enterprise AI capabilities and synergies across the Nebius ecosystem.
- Efficiency improvement and cost reductions have improved segment financial performance.

Equity stakes

In addition to our non-core businesses, we own equity stakes in both Toloka and ClickHouse, both of which were originally spun out from Nebius Group.

Toloka

Toloka is a leading provider of scalable, high-quality data essential for LLM, GenAI, and AI agent development.

To capture value across the entire AI development stack, Toloka recently executed several strategic platform expansions:

- **Toloka Arena:** An independent suite of benchmarks that evaluates agentic intelligence on private, non-contaminated tasks to simulate real-world scenarios involving multi-turn tool use, live databases, and strict business rules. With over 40 models evaluated, Arena empowers enterprises to make confident model selections and enables AI labs to license these proprietary datasets to train and test their own models.
- **Self-service platform:** A major upgrade to the data-labelling platform, accelerating time-to-scale for data pipelines through API automation, agent-assisted setup, built-in quality controls, and newly integrated synthetic data generation.
- **Expanded catalog:** A significantly broader catalog of expert-validated datasets across high-demand domains, including coding, STEM, reasoning, and RL Gyms.
- **Physical AI and robotics:** Accelerated investment in physical AI, expanding robotics data offerings beyond egocentric video to include UMI-style data and teleoperation.

As of Q2'25 we no longer hold voting control of Toloka. However, we maintain a significant equity stake, and are encouraged by the growing investor interest in the AI data provider market.

ClickHouse

ClickHouse is an open-source database management system built for real-time data processing and analytics.

In January 2026, it was reported that ClickHouse raised $400M in a Series D financing at a valuation of approximately $15B. Following this financing, Nebius Group continues to own a significant minority equity stake in ClickHouse.

Financial update

Revenue

In USD $ millions	Three months ended June 30			Six months ended June 30		
	2025	**2026**	**Change**	**2025**	**2026**	**Change**
Revenues	105.1	582.3	*454%*	156.0	981.3	*529%*

Nebius Group once again executed against its financial goals. Q2 group revenue of $582.3 million, was up 454% year-over-year, and up 46% compared to Q1. Year-over-year growth was driven by capacity scaling and supported by higher pricing and improved utilization from a year ago.

- Nebius AI cloud revenue was $574.9 million in Q2'26, a 514% increase year-over-year from Q2'25. Our Nebius AI cloud business accounted for approximately 98% of total group revenue during the quarter.
- Annualized run-rate revenue (ARR)1 of $3.0 billion as of the end of June was up 598% year-over-year and 56% from the $1.9 billion reported at the end of March.

Operating expense

In USD $ millions	Three months ended June 30			Six months ended June 30		
Expense category	**2025**	**2026**	**Change**	**2025**	**2026**	**Change**
Cost of revenues	30.1	133.6	*344%*	54.8	237.4	*333%*
as a percentage of revenues	29%	23%		35%	24%	
Product development	42.8	191.0	*346%*	79.3	258.4	*226%*
as a percentage of revenues	41%	33%		51%	26%	
Sales, general and administrative	68.2	173.9	*155%*	129.1	317.7	*146%*
as a percentage of revenues	65%	30%		83%	32%	
Depreciation and amortization	75.2	259.7	*245%*	124.3	471.7	*279%*
as a percentage of revenues	72%	45%		80%	48%	
Total operating costs and expenses	**216.3**	**758.2**	*251%*	**387.5**	**1,285.2**	***232%***
as a percentage of revenues	206%	130%		248%	131%	

Cost of revenue was $133.6 million in Q2'26, representing an increase of 344% compared to the same period in 2025.

- The increase was due to the expansion of our Nebius AI cloud business, with expenses incurred for co-location and operating lease agreements as well as hiring to support our growing operations.
- As a percentage of revenue, cost of revenue was 23% in Q2'26, down from 29% in Q2'25, primarily reflecting operating leverage as we scaled capacity.

Product development expenses were $191.0 million in Q2'26, representing an increase of 346% compared to the same period in 2025.

- The increase in product development expenses was primarily driven by hiring in our engineering and development teams to build and enhance our product offerings, as well as $115.9 million of non-recurring expense pertaining to share-based compensation expense and post-combination expenses recognized in connection with the acquisition of Eigen AI.
- As a percentage of revenue, product development expenses decreased to 33% from 41% in the prior year.

Sales, general and administrative expenses ("SG&A") were $173.9 million in Q2'26, representing an increase of 155% compared to $68.2 million in the same period in 2025.

- The increase was primarily driven by hiring to support the growth of our business as well as consulting, legal and other professional fees primarily incurred in connection with acquisition-related activities and financing transactions.
- SG&A as a percentage of revenues declined to 30% from 65% in Q2'25.

Depreciation and amortization expenses ("D&A") were $259.7 million in Q2'26, representing an increase of 245% compared to the same period in 2025.

- The primary driver of the dollar increase in D&A expenses was the continued investments in GPU-related capital expenditures and related data center hardware for the Nebius AI cloud business.
- We use a five-year useful life for our server and network equipment based on usage patterns and current utilization commitments, up from the four-year useful life used prior to 2026.
- D&A as a percentage of revenue declined to 45% from 72%.

Adjusted EBITDA

In USD $ millions	Three months ended June 30			Six months ended June 30		
	2025	**2026**	**Change**	**2025**	**2026**	**Change**
Adjusted EBITDA / (loss)	(21.0)	236.2	n/m	(74.7)	365.7	*n/m*
as a percentage of revenues	-20%	41%		-48%	37%	

Group adjusted EBITDA of $236.2 million in Q2 improved by $257.2 million year over year.

- The continued improvement in Adjusted EBITDA was driven primarily by the strong growth in our Nebius AI cloud business, which generated adjusted EBITDA of $285.7 million and adjusted EBITDA margin of 49.7%.
- We improved group adjusted EBITDA from Q1'26 by $106.7 million from $129.5 million to $236.2 million.

Capital expenditures

In Q2'26, capital expenditures were approximately $5.7 billion, primarily driven by purchases of GPUs and GPU-related hardware, and our data center expansion activities.

Capital requirements

We will continue to invest in capex throughout the year, and intend to use a diversified range of funding sources including:

Cash generated from operations and upfront customer payments.

- We expect to receive over $9 billion in customer prepayments in 2026.
- Upfront payments give us visibility into future cash flows and flexibility on timing of additional financing.

We may also enter into debt financing transactions or access the debt capital markets.

- We are actively progressing potential debt transactions, including asset-backed financing and corporate-level debt.
- In July we entered into our first secured debt financing, raising $775 million priced at SOFR + 2.50%. The vehicle is backed by deployed GPU infrastructure and contracted cash flows from an agreement with an investment-grade customer and demonstrates our ability to fund growth on attractive terms. We have more than $40 billion of customer commitments.
- We plan to continue tapping into these financing options.

Our financing options include our at-the-market (ATM) program, which we began using this quarter.

- Through June 30, 2026, we sold 12.7 million Class A shares at a weighted-average price of $223.6 per share, generating gross proceeds of approximately $2.8 billion.
- As of June 30, 12.3 million shares remained available under the program.

And we may also evaluate other financing options, and will ultimately pursue the vehicles that best serve the long-term interests of the business.

Guidance

The company will share a detailed view of guidance on its earnings call and webcast.

Earnings webcast

Nebius Group will host a conference call and earnings webcast at 5:00 a.m. Pacific time/8:00 a.m. Eastern time/2:00 p.m. Central European time on August 12, 2026, to discuss these financial results. To register to participate in the conference call, or to listen to the live audio webcast, please visit Nebius's Investor Relations website at group.nebius.com/investor-hub.

A replay will be available on the same website following the call.

Forward-looking statements

This document contains forward-looking statements that involve risks and uncertainties. All statements contained or implied other than statements of historical facts, including, without limitation, statements regarding our business plans, market opportunities, capacity buildout plans, capital expenditure requirements, financing requirements and projected financial performance, are forward-looking statements. In some cases, these forward-looking statements can be identified by words or phrases such as "may," "will," "expect," "anticipate," "aim," "estimate," "intend," "plan," "believe," "potential," "continue," "is/are likely to" or other similar expressions. In addition, these forward-looking statements reflect our current views with respect to future events and are not a guarantee of future performance. Actual results may differ materially from the results predicted or implied by such statements, and our reported results should not be considered as an indication of future performance. The potential risks and uncertainties that could cause actual results to differ from the results predicted or implied by such statements include our ability to: obtain sufficient financing and manage our liquidity and capital resources to support our operations and growth; successfully identify, develop and bring online additional data center capacity on a timely and cost-effective basis, including securing suitable sites and access to power; implement and maintain effective internal control over financial reporting; manage supply chain risks and secure required equipment, hardware, materials and services on acceptable terms; compete effectively in a dynamic and competitive market while generating sustained customer demand; and manage dependence on key vendors and adapt to technological change.

Many of these risks and uncertainties depend on the actions of third parties and are largely outside of our control. Our actual results of operations may also differ materially from those stated in or implied by such forward-looking statements as a result of a variety of factors, including those described under the captions "Risk Factors" and "Operating and Financial Review and Prospects" in our Annual Report on Form 20-F for the year ended December 31, 2025 filed with the U.S. Securities and Exchange Commission ("SEC") on April 30, 2026, which is available on our investor relations website at https://group.nebius.com and on the SEC website at www.sec.gov. All information in this document is as of the date hereof, and the Company undertakes no duty to update this information unless required by law.

In addition, statements that "we believe" and similar statements reflect our beliefs and opinions on the relevant subject. These statements are based upon information available to us as of the date of this document, and while we believe such information forms a reasonable basis for such statements, such information may be limited or incomplete, and our statements should not be read to indicate that we have conducted an exhaustive inquiry into, or review of, all potentially available relevant information. These statements are inherently uncertain, and investors are cautioned not to unduly rely upon these statements.

We operate in an evolving environment. New risks emerge from time to time, and it is not possible for our management to predict all risks, nor can we assess the effect of all factors on our business or the extent to which any factor, or combination of factors, may cause actual results to differ materially from those contained in any forward looking statements. You should not rely upon forward looking statements as predictions of future events. We undertake no obligation to update or revise any forward-looking statements, whether as a result of new information, future events or otherwise.

Disclaimer

Links to third-party websites are provided for informational purposes only; Nebius is not responsible for the content contained on or accessible through the linked sites.

Use of Non-GAAP financial measures

To supplement the financial information prepared and presented in accordance with U.S. GAAP, we present the following non-GAAP financial measures: Adjusted EBITDA/(loss) and Adjusted net income/(loss). The presentation of these financial measures is not intended to be considered in isolation or as a substitute for, or superior to, the financial information prepared and presented in accordance with U.S. GAAP. For more information on these non-GAAP financial measures, please see the tables captioned "Reconciliations of non-GAAP financial measures to the most directly comparable U.S. GAAP measures", included following the accompanying financial tables. We define the various non-GAAP financial measures we use as follows:

- Adjusted EBITDA / (loss) means U.S. GAAP net income/(loss) from continuing operations before (1) depreciation and amortization, (2) share-based compensation expense, (3) acquisition and other corporate transaction-related costs, (4) interest income, (5) interest expense, (6) income/(loss) from equity method investments, (7) gain from revaluation of investments in equity securities, (8) other income/(loss), net, (9) income tax expense/(benefit).
- Adjusted net income / (loss) means U.S. GAAP net income/(loss) from continuing operations before (1) share-based compensation expense, (2) acquisition and other corporate transaction-related costs, (3) amortization of debt discount and issuance costs, net of interest expense capitalized, (4) foreign exchange gains/(losses) and (5) gain from revaluation of investments in equity securities. Tax effects related to the listed adjustments are excluded from adjusted net income.

These non-GAAP financial measures are used by management for evaluating financial performance as well as decision-making. Management believes that these metrics reflect the organic, core operating performance of the company, and therefore are useful to analysts and investors in providing supplemental information that helps them understand, model and forecast the evolution of our operating business.

Although our management uses these non-GAAP financial measures for operational decision-making and considers these financial measures to be useful for analysts and investors, we recognize that there are a number of limitations related to such measures. In particular, it should be noted that several of these measures exclude some recurring costs, particularly share-based compensation. In addition, the components of the costs that we exclude in our calculation of the measures described above may differ from the components that our peer companies exclude when they report their results of operations.

Below we describe why we make particular adjustments to certain U.S. GAAP financial measures:

Net income/(loss) from discontinued operations

We present Adjusted EBITDA/(loss) and Adjusted net income/(loss) excluding any effects of our discontinued operations.

Information on our discontinued operations is disclosed in our Annual Report on Form 20-F for the year ended December 31, 2025 filed with the U.S. Securities and Exchange Commission ("SEC") on April 30, 2026.

Share-based compensation expense

Share-based compensation is a significant expense item and an important part of our compensation and incentive programs. As it is highly dependent on our share price at the time of equity award grants, we believe that it is useful for investors and analysts to see certain financial measures excluding the impact of these charges in order to obtain a clearer picture of our operating performance.

Foreign exchange gains/(losses)

The functional currency of Nebius Group N.V. is the United States Dollar, which is also the Group's reporting currency. Foreign exchange gain/(loss) dynamics reflect changes in the U.S. dollar value of monetary assets and liabilities that are denominated in other currencies, as well as changes in the functional currencies of foreign subsidiaries' monetary assets and liabilities that are denominated in currencies different from their respective local currencies. Because foreign exchange fluctuations are outside of our operational control, we believe that it is useful to present Adjusted EBITDA/(loss), adjusted net income/(loss) and related margin measures excluding these effects, in order to provide greater clarity regarding our operating performance.

Acquisition and other corporate transaction-related costs

We believe that it is useful to present Adjusted net income/(loss), Adjusted EBITDA/(loss) and related margin measures excluding impacts not related to our operating activities. Adjusted net income/(loss) and Adjusted EBITDA/(loss) exclude certain expenses related to M&A activities and other expenses related to corporate transactions.

Amortization of debt discount and issuance costs, net of interest expense capitalized

We also adjust net income/(loss) for interest expense representing amortization of the debt discount and issuance costs related to our convertible senior notes, net of interest expense capitalized into cost of our property and equipment. Debt discount represents the accretion of the nominal amount of notes payable at maturity, unless the relevant notes have been earlier repurchased, redeemed or converted in accordance with their terms. We adjust net income/(loss) for the interest expense recognized from amortization of the debt discount and issuance costs due to the significantly different timing of payment in relation to the operating results.

Contact investor relations:
askIR@nebius.com