


Nebius Group Letter to shareholders Q1 2026

Raising capacity guidance for YE'26



Record pipeline generation in Q1

Growing demand drove record pipeline generation, up ~3.5x QoQ in Q1

Foundational partnerships

Meta

Favorable terms enable buildout of the Nebius cloud ecosystem

Largest deal in our history

\$27B

NVIDIA

\$2B investment

Expanded partnership with deeper technical collaboration, powering inference and agentic AI

Product — Strategic acquisitions deepen competitive differentiation

Eigen AI: Inference at scale

Announced May 2026

Fast AI inference — ranked #1 by NVIDIA at GTC 2026

High throughput per GPU lowers cost per token

Improves margins on GPU fleets and scales efficiently

Strengthens Token Factory platform and brings industry leading talent to Nebius

Clarifai: Inference optimization

Announced May 2026

System-level GPU optimization cuts costs for complex open-source models

Reduces compute capacity needed to run large models in production

Nebius AI Cloud "Aether 3.5"

Serverless AI delivers instant on-demand inference without capacity planning

Data transfer removes friction from enterprise migration

Strengthens security, governance, and operational controls

Unlocks physical AI workloads

Scaling capacity — Global AI cloud

New owned sites:

1.2GW AI factory in PA

300MW+ site in Finland

of contracted power

Number of sites with >100MW

YE'25: 1 | Current: 7

Q1 financial highlights

2026 Guidance update

On track to achieve \$3.0B-\$3.4B revenue in 2026 and \$7B-\$9B ARR

Q4'2025: \$1.25B | Q1'2026: \$1.9B | Guidance FY'26: \$7B-\$9B

Q1 Nebius AI cloud revenue

Up 841% YoY and 82% QoQ to \$390M

YoY and QoQ growth rates accelerated from Q4

Significant Q1 fundraising

\$6.3B capital secured

In convertible notes and an equity investment from NVIDIA

On track to achieve ~40% Adj. EBITDA margin in 2026

Strong cash position ending Q1

\$9.3B cash

Including \$2.3 billion in positive operating cash flow in Q1

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

More than 2x size of autonomous fleet YTD

Strategic investment from Uber

tripleten

Revenue growth of 10% YoY in Q1 driven by strong customer acquisition and expanding tech program offerings

Note: Annualized run-rate revenue (ARR) is calculated by taking Nebius AI cloud revenue from the last month of the quarter multiplied by 12

Dear shareholders,

We continue to see unprecedented demand across the market. Compute and cloud needs are vastly exceeding capacity as more industries embrace AI and companies move beyond experimentation to real-world applications. We are seeing this demand first hand, and are capturing it with our full-stack AI-native cloud.

We are a technology company at our core. We have world-class engineers and deep proprietary expertise across every layer of the stack. From infrastructure and multi-tenant cloud through inference solutions and agentic platforms. We are not simply responding to where the industry stands today; we have the knowledge and experience to build the infrastructure, tools, and capabilities for where it will be tomorrow.

The results of the first quarter are on the right path.

- **Consistent execution drives strong financial results**

ARR¹ grew 674% year-over-year, while revenue accelerated both quarter-on-quarter and year-on-year. Our deepening customer relationships are strengthening our unique position to serve industry-wide demand for both compute and cloud services.

We also demonstrated the operating leverage inherent in our business. Nebius, our AI cloud business, nearly doubled its adjusted EBITDA margin quarter-on-quarter to 45%, a testament to our trajectory toward 20-30% EBIT margins.

- **Strategic acquisitions deepen competitive differentiation**

We announced three acquisitions that **advance our capabilities** beyond the infrastructure layer of the stack and **into inference and agentic workloads**:

In February, we announced the acquisition of Tavily to expand our platform with market-leading agentic search capabilities;

Earlier this month, we announced an agreement to acquire Eigen AI to strengthen Nebius Token Factory as a frontier managed inference platform;

Yesterday, we also announced the acquisition of Clarifai brings in significant new engineering and research talent, as well as what we believe is the company's leading inference and compute orchestration technology.

These acquisitions substantially strengthen our ability to serve new segments in the market as well as an expanded set of customers across the AI workload lifecycle.

They also show **Nebius is a top destination for industry-leading talent, and establish the company's engineering presence in the Bay Area.**

The release of Nebius AI Cloud Aether 3.5 in March made our hyperscale-grade platform more frictionless, removing barriers between ideas and real-world AI solutions by providing compute that is powerful, efficient, and built to scale.



Arkady Volozh,
Founder and CEO

“We are not simply responding to where the industry stands today; we have the knowledge and experience to build the infrastructure, tools, and capabilities for where it will be tomorrow.”

- **Go-to-market remains agile as the market evolves**

Pipeline generated in the quarter reached a new record, increasing ~3.5x quarter-on-quarter.

We remain agile on capacity allocation and contract terms. We proactively manage our ability to support on-demand capacity while also offering reserve contracts to give customers flexibility based on their needs.

We are investing across the company to ensure we have the right team in place to capitalize on the full opportunity ahead. One such area is customer success, which will fuel our land-and-expand model to support our healthy expansion rate across use cases and workloads.

- **Owned capacity continues to expand rapidly**

Contracted capacity already exceeds 3.5 GW, far surpassing the goal of 3 GW we set for the end of the year.

Our execution gives us the confidence to raise our contracted power guidance to more than 4 GW by year-end.

We have now secured two gigawatt-scale sites in the US. We broke ground at our Missouri location yesterday; today, **we are announcing a new site in Pennsylvania** where we have secured power for a **deployment of up to 1.2 GW.**

Owned capacity now accounts for more than 75% of our contracted power.

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

We remain laser-focused on bringing this contracted capacity online. **We expect capacity added in Q3 to significantly expand our footprint.**

- **Key partnerships continue to support growth**

We landed a second large agreement with Meta for up to \$27 billion, giving us substantial new capital to accelerate the growth of our AI cloud business and pursue potential additional financing at attractive rates.

We **achieved NVIDIA Exemplar Cloud status on GB300 NVL72 for training**. Nebius is among the first cloud providers globally to achieve this designation, and one of a small group to hold it across multiple GPU generations.

We deepened our partnership with NVIDIA, expanding our **software and hardware integration** including AI factory design and support, and collaboration on building out a best-in-class inference and agentic stack around Nebius Token Factory.

NVIDIA also made a **\$2 billion investment in Nebius**, reflecting their continued confidence in our business and unique depth of engineering expertise across the full AI tech stack.

- **Successful fundraising supports disciplined capacity expansion**

In Q1, **we raised \$6.3 billion**, including the \$2 billion equity investment from NVIDIA, and \$4.3 billion from convertible securities at favourable rates.

Our balance sheet is healthy, with more than \$9 billion of cash from fundraising during the quarter and robust Q1 operating cash inflows of \$2.3 billion. We plan to deploy this cash in a disciplined manner to further expand our capacity, and will evaluate potential acquisitions that deepen our technology stack and product capabilities.

The achievements that our team delivered in the first quarter set us up for continued success. Our capacity footprint is expanding rapidly, our full-stack cloud platform is world-class from the infrastructure layer all the way up to our inference and agentic capabilities.

We will continue to execute, while maintaining a relentless focus on serving our customers around the world.

The future is very bright, and we have a long and exciting journey ahead.

Arkady Volozh

Group Q1 update



NEBIUS

Building the AI cloud for production at scale

In Q1, we made significant strides in expanding our platform toward a more complete full-stack. Through organic product development and targeted acquisitions, we strengthened our capabilities across the entire AI lifecycle to best offer a optimized cloud for production AI at scale.

This will allow us to bring a broader set of users into the ecosystem, such as:

- Data researchers accelerating model training and experimentation,
- AI engineers deploying and scaling inference without managing infrastructure complexity,
- Enterprise teams — from IT to DevOps — operating with the governance, security, and operational control.

Our latest Aether 3.5 release enhances enterprise readiness and introduces serverless AI capabilities. Key features include:

- Serverless, allowing customers to experiment, optimize models, and deploy workloads more efficiently within a flexible and predictable environment.
- A new Data Transfer Service that enables data movement across clouds at petabyte scale, giving customers greater flexibility in how they manage and migrate data.
- Strengthened security, governance, and operational controls through improved Kubernetes Secrets integration, as well as expanded billing and audit log export capabilities.

Strengthening the stack through targeted acquisitions

Nebius Token Factory, our inference platform that enables customers to deploy and optimize open-source and custom models, was released last quarter and is already seeing momentum — across new and existing customers as it reduces operational overhead.

In addition to organic, in-house development, we are continuing to strengthen the software layer through strategic acquisitions. Since the beginning of the year we have added advanced inference optimization and agentic AI capabilities through our:

- Acquisition of Tavily, which will expand retrieval and agentic workflows;
- Announced an agreement to acquire Eigen AI to enhance inference execution and model optimization; and
- Welcomed Clarifai's team alongside an IP license agreement to further strengthen system-level inference optimization.

Eigen AI strengthens Nebius Token Factory as a managed inference platform for production AI, and adds leading inference research and post-training optimization to our in-house R&D organization. Our teams have already delivered jointly-optimized endpoints that achieved top rankings on Artificial Analysis across multiple models.

Where Eigen AI's focus is on model optimization, Clarifai operates at the system level, building the end-to-end infrastructure required to run complex open-source models reliably in production. Clarifai's founder and CEO Matthew Zeiler — a recognized pioneer in machine learning — will join Nebius as SVP of Research, leading a dedicated unit focused on frontier AI innovation across areas including multimodal agentic reasoning, world models, token efficiency, and long-term memory.

Alongside these acquisitions, we are deepening our engineering collaboration with NVIDIA on agentic and inference software — work we believe can become foundational across the industry. We formalized this evolution of our partnership in March. Together, these advance our vision of an integrated AI platform spanning infrastructure, training, inference, retrieval, optimization and agentic capabilities.

Delivering efficient, AI-optimized infrastructure at scale

Our infrastructure is purpose-built to help customers train, fine-tune, and deploy models quickly and cost-effectively.

In Q1, we **achieved NVIDIA Exemplar Cloud status on GB300 NVL72 for training workloads**, reinforcing the performance of our platform. Nebius is among a limited group of providers to achieve this status across multiple GPU generations.

We also broadened our platform to include NVIDIA RTX PRO 6000 Blackwell Server Edition on Aether 3.5, enabling customers to run a wider range of applied AI and simulation workloads.

With **NVIDIA Vera Rubin NVL72 arriving from the second half of 2026** and gigawatt-scale AI factories in the active construction phase, Nebius continues to strengthen the ways it enables customers to build and scale AI faster than ever before.

Scaling our infrastructure globally

In this section, we discuss three forms of power:

- (1) **contracted power**, or power secured by contracted land and power commitments;
- (2) **connected power**, or power connected into fully built and equipped data centers; and
- (3) **active power**, or power being consumed by installed, operational IT equipment and available for revenue generation.

The first quarter was once again defined by strong execution and significant progress toward our capacity expansion targets.

We are pursuing a twin-track strategy: bringing capacity online to support near-term growth, and securing land and power commitments that extend well beyond this year.

Capacity underpinning 2026 growth

In Q1, **we delivered on all capacity commitments** across our AI cloud customers and strategic long-term contracts (Microsoft, Meta). We remain focused on bringing additional capacity online this year, and are progressing toward delivery across both new and existing co-location sites. We expect to significantly increase our capacity in place in 3Q.

We continue to expect 800MW to 1GW of connected power by year-end.

Building the foundation for future growth

Today we are **announcing our new second owned gigawatt-scale site in the United States — in Pennsylvania**, with up to 1.2 GW of power. The planned AI factory at this location will be delivered in phases beginning in 2027 and will add to Nebius's rapidly expanding US footprint, alongside the 1.2 GW AI factory in Independence, Missouri.

In the quarter we also announced a new owned location in Finland that when fully deployed at 310 MW will be one of Europe's largest dedicated AI factories.

This led us to surpass our 2026 year-end contracted power target of 3 GW.

Our **contracted capacity now exceeds 3.5 GW**, with owned capacity representing more than 75% of the total.

Given this momentum, we are raising our contracted power guidance to more than 4 GW by year-end, with any additional contracts contributing to capacity growth in 2027 and beyond.

In addition to the new owned sites in Pennsylvania and Finland, we are continuing to build out our global footprint:

- We are in the **active construction phase at several of our owned sites**, including Missouri and Alabama, which we expect to be operational in 2027.
- In New Jersey, we continue to service our customer commitments and remain on track to activate the remaining capacity through the year, with the majority coming online in the second half of the year.
- We also added a new colocation site in Spain.

In total, our owned facilities will deliver 3 GW of capacity across five sites, reinforcing our commitment to a capital-efficient model with attractive long-term unit economics.

Beyond scale, we continue to lead on technology. Nebius will be among the first AI cloud providers worldwide to deploy NVIDIA Vera Rubin NVL72 systems, with deployments planned across both the US and EMEA.

Expanding and diversifying our customer base

Our ability to deliver both high-performance, large-scale clusters and smaller-scale, on-demand compute gives us a distinct advantage. In Q1, that advantage translated into a sequential re-acceleration in revenue.

The momentum is broad-based and compounding:

We saw **record pipeline generation** and sequential revenue re-acceleration in the quarter, with **our pipeline up approximately 3.5x quarter over quarter**.

Pricing continued to rise for new generation GPUs, with older generation chips also seeing strong pricing support;

Average deal sizes also grew across both new and existing customers, driven by pricing, GPU commitments, and longer duration.

Critically, this growth is coming from a widening base of customers across verticals, model builders, and inference workloads.

Vertical GTM is winning customers across high-growth segments

Our vertical go-to-market approach is delivering wins across key customer segments.

Industry spotlight: healthcare and life sciences

Key customer relationships this quarter include:

- **Sword Health**, which selected Nebius's AI cloud to power Dawn, its new direct-to-consumer mental health and wellbeing solution, and Thrive, its AI-guided care platform for pain and musculoskeletal recovery (MSK). Our platform seamlessly scaled Dawn's text-based model to 200+ billion parameters, while supporting Thrive's vision AI models for movement analysis, all while maintaining excellent compute performance, service, and reliability.
- A leader in generative AI for molecular biology achieved over 4x training speedups and faster inference on Nebius's full-stack AI infrastructure — performance gains that translate directly into faster research cycles and accelerated drug discovery.
- AI-native startups across the vertical are choosing Nebius to:
 - Diagnose and treat complex diseases;
 - Build a CAD suite for molecules to make drug discovery faster, cheaper, and more precise;
 - Train a 3B+ parameter model fine-tuned on proprietary cancer datasets to discover drugs for patients who haven't responded to existing therapies.
 - Speed scientific breakthroughs through accelerated virtual experimentation, achieving nucleotide-level tokenization and a 100%+ increase in the accuracy of paired mRNA and ribosome predictions;

Physical AI

Startups in **physical AI and robotics** need large-scale GPU training, simulation testing, and edge deployment, and we are already playing a pivotal role propelling this technology forward with industry leaders. In March, **we announced an agreement to collaborate with NVIDIA to accelerate physical AI development** with an end-to-end platform built for the full robotics lifecycle, from simulation and training to real-world deployment. That leadership is translating into significant recent wins including:

In Q1, we signed an agreement with **1X Technologies**, a leading robotics company building general-purpose robots capable of performing any kind of work autonomously.

We recently added **Rhoda** to our customer ecosystem, and will enable its work building robotic intelligence based on huge multi-modal data sets to deliver video-predictive control to help robots react to and learn from physical world operating environments.

To capitalize on the opportunity we see in robotics and physical AI, we introduced a physical AI solution that combines synthetic data generation, orchestration, world models, and real-world inference into a managed platform offering. This enables enterprise customers and design partners to develop and deploy

vertical AI applications more efficiently across emerging real-world use cases.

We continue to power leading AI natives and model builders. In Q1, we signed a number of new deals with customers including:

- **Core Automation**, which selected Nebius AI cloud to accelerate its effort to rethink neural network architecture and reinvent how foundational models are built, and
- **Logical Intelligence**, a next-gen model builder using the Nebius platform to develop AI systems for reasoning that move beyond statistical probability to mathematical certainty.

Nebius Token Factory

Open-source models are rapidly improving and winning adoption as high-growth startups move from experimentation to production at scale.

Significant Token Factory customer wins in Q1 include:

- **Revolut**, which is developing a platform to simplify finance for businesses and consumers, was able to remove human intervention from 80% of support chats, and handle 1.2M chat tickets per month on Token Factory; and
- **monday.com**, which selected Token Factory to support its AI work platform that helps manage, orchestrate, and execute workflows.

Building the global sales and GTM organization

In March, we announced two important leadership hires. Dan Lawrence was appointed as SVP and GM for the Americas. Dan brings deep experience building and operating cloud businesses at multi-billion-dollar scale. Most recently, he served as SVP of Global Sales for Cloud at Akamai Technologies, where he built the go-to-market model and rapidly scaled its compute business. Prior to Akamai, he held senior leadership roles at Amazon Web Services.

John Haarer was appointed as GM for Asia-Pacific and Japan. John brings over a decade of experience driving go-to-market efforts in the region for global technology leaders including Cloudflare and Twilio.

We have also welcomed Raja Agrawal as VP of Sales for the Middle East and Africa, based in Dubai. Raja brings more than two decades of leadership experience at the intersection of enterprises and cloud, with deep expertise in AI, cloud, data, and enterprise platforms from senior leadership roles at SAP, Microsoft and Browserstack.

Across the GTM organisation, we are also growing our sales reps and customer success teams to support our global expansion.

With senior leadership established across key regions and our compute capacity and sales team scaling rapidly, we are well

positioned to capture the significant wave of demand as AI capabilities become foundational to the global economy.

Strategic agreements

Strategic long-term contracts remain an important part of our financing strategy.

Our latest **\$27 billion, five-year agreement with Meta**, is uniquely structured and provides important benefits.

The deal comprises two parts:

- A **\$12 billion**, five-year purchase of compute capacity that is scheduled to begin in early 2027;
- An additional **\$15 billion** contract that allows Nebius to sell capacity to Meta on pre-agreed terms, or to our AI cloud customers at market rates.

This structure gives us attractive financing options together with **long-term revenue visibility, and potential upside** as Nebius can allocate capacity to Meta throughout the term of the contract.

Avride

Autonomous vehicle operations delivered significant achievements in Q1

- The team has already more than doubled the size of the AV-capable vehicle fleet vs. year-end 2025 levels.
- Expansion of their fleet and operating map in Dallas are expected to continue throughout the year.
- Avride's streamlined process for pre-assembling key electronic components, including the rooftop sensor suite, has driven higher retrofitting throughput and accelerated fleet deployment.
- The company is scaling its R&D fleet to provide the high-fidelity data required to refine its state-of-the-art AI stack and accelerate the transition to No-Vehicle-Operator (NVO) operations.

Robodelivery operations launched in Philadelphia and campus initiatives advanced, with new locations and expected launches planned later this year

- Robot deliveries increased 178% YoY in Q1 to over 174,000 for the quarter, and surpassed 500,000 since inception in early April.
- In Q1, Avride launched robodelivery operations in Philadelphia - the first delivery robots in the city. This expansion represents a new milestone of the company's multi-year partnership with Uber to bring Avride's delivery robots and autonomous vehicles to Uber and Uber Eats customers across the US.
- Robot deliveries began at Salisbury College during Q1, achieving an on-schedule launch, solid early utilization

rates, and initial delivery results that exceeded customer expectations.

- The FY26 growth plan is supported by a robust pipeline of new campuses and additional cities, each representing a natural and replicable extension of our proven operating model.

TripleTen

TripleTen, our edtech business, continues to develop its offerings to meet shifting user demand

- TripleTen is meeting demand for AI skills by integrating additional AI concepts across its web development, software development, and Quality Assurance programs.
- The company launched additional AI upskilling programs within the B2B platform during the first quarter.
- Product development tested new solutions aimed at supporting less experienced candidates navigating a challenging entry-level job market.
- TripleTen revenue growth was 10% year over year in Q1.
- Regionally, Latin America and Brazil markets continued to post stable growth, outperforming the US market.
- An increased emphasis on efficiency and profitability is expected to support future financial results.
- New student growth in Q1 was approximately 5,000.

Equity stakes: Toloka and ClickHouse

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 data provider for LLM and GenAI developers, delivering scalable, high-quality, curated data for AI agents and model development. The company serves major frontier model producers, as well as other leading frontier labs, hyperscalers and technology enterprises.

In February, we announced plans to integrate Toloka's Tendem solution into the Nebius ecosystem. Originally designed as a hybrid human-AI agent, Tendem was the first platform to embed vetted human experts directly into agentic workflows — making expert judgment callable via the Model Context Protocol (MCP), the emerging standard for AI tool integration. This integration further strengthens the Nebius AI stack, anchoring the raw intelligence of Token Factory and the autonomy of Tavily's agentic search with a programmable layer of human reliability.

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. The re-valuation of Nebius Group's equity stake following this financing contributed a gain of \$781M to non-operating income in the first quarter. This is a non-cash item that captures the growth in the underlying value of our stake.

Following this financing, Nebius Group continues to own a significant minority equity stake in ClickHouse.

Financial update

Revenue

In USD \$ millions

Three months ended March 31

	2025	2026	Change
Revenues	50.9	399.0	684%

Nebius Group once again executed against its financial goals. Q1 group revenue of \$399.0 million exceeded our expectations, up 684% year-over-year, and up 75% compared to Q4. Growth was driven by capacity scaling and supported by strong pricing and utilization.

- Nebius AI cloud revenue was \$389.7 million in Q1'26, an 841% increase year-over-year from Q1'25. Our Nebius AI cloud busi-

ness accounted for approximately 98% of total group revenue during the quarter.

- Annualized run-rate revenue (ARR)¹ of \$1.92 billion as of the end of March was up 674% year-over-year and 54% from the \$1.25 billion reported as of the end of December 2025.

Operating expense

In USD \$ millions

Three months ended March 31

Expense category	2025	2026	Change
Cost of revenues	24.7	103.8	320%
as a percentage of revenues	49%	26%	
Product development	36.5	67.4	85%
as a percentage of revenues	72%	17%	
Sales, general and administrative	60.9	143.8	136%
as a percentage of revenues	120%	36%	
Depreciation and amortization	49.1	212.0	332%
as a percentage of revenues	96%	53%	
Total operating costs and expenses	171.2	527.0	208%
as a percentage of revenues	336%	132%	

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

Cost of revenue was \$103.8 million in Q1'26, representing an increase of 320% 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 26% in Q1'26, down from 49% in Q1'25, primarily reflecting operating leverage as we scaled capacity.

Product development expense were \$67.4 million in Q1'26, representing an increase of 85% 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 a percentage of revenue, product development expenses decreased to 17% from 72% in the prior year.

Sales, general and administrative expenses ("SG&A") were \$143.8 million in Q1'26, representing an increase of 136% com-

pared to \$60.9 million in the same period in 2025.

- The increase was primarily driven by hiring to support the growth of our business.
- SG&A as a percentage of revenues declined to 36% from 120% in Q1'25.

Depreciation and amortization expenses ("D&A") were \$212.0 million in Q1'26, representing an increase of 332% 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.
- Starting Q1'26, we revised the useful life for our server and network equipment from four years to five years to reflect usage patterns and current utilization commitments. The change in accounting estimate has been applied prospectively from 2026.
- D&A as a percentage of revenue declined to 53% from 96%.

Adjusted EBITDA

In USD \$ millions

Three months ended May 31

	2025	2026	Change
Adjusted EBITDA / (loss)	(53.7)	129.5	n/m
as a percentage of revenues	-106%	32%	

Group adjusted EBITDA of \$129.5 million in Q1 improved by \$183.2 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 \$174.0 million and adjusted EBITDA margin of 45%.

- We improved group adjusted EBITDA from Q4'25 by \$114.5 million from \$15.0 million to \$129.5 million.

Capital expenditures

In Q1'26, capital expenditures were approximately \$2.5 billion, primarily driven by purchases of GPUs and GPU-related hardware, and our data center expansion activities.

- Debt: We may from time to time seek to 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. We plan to start tapping into these financing options by raising mid-single digits billions of dollars in the near term.
- On top of that, our financing options include our at-the-market program. We have not utilized this program to date, but are evaluating the program regularly.
- And we are very focused on generating prepayments from our current and future customers, in order to reduce the capital needed from equity and debt financing.
- We may also evaluate other financing options, but will ultimately pursue whichever vehicles best serve the long-term interests of the business.

Capital requirements

We will continue to invest in capital expenditures throughout the year. We will leverage a diversified range of funding sources. This includes:

- Cash we generate from operations and upfront customer payments.

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 May 13, 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 nearest 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) SBC expense, (3) one-off restructuring and other expenses, (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) SBC expense, (2) one-off restructuring and other expenses, (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.

SBC expense

SBC (Stock-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.

One-off restructuring and other expenses

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 the restructuring, M&A activities and other similar one-off expenses.

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

