

# Nebius Group Letter to shareholders Q3 2025

## Strong momentum into 2026

# \$7–9 billion

expected annualized run-rate revenue\* at year-end 2026

\*Calculated by taking revenue from the last month of the quarter multiplied by 12

## Significantly raised capacity guidance

Contracted power at CYE 2026

1 GW

Previously announced plan

>2.5 GW

Current plan

Connected power

Includes power that is fully provisioned and can be activated immediately upon GPU installation

800 MW–1 GW

25 MW

CYE 2024

220 MW

CYE 2025

CYE 2026

## Q3 financial highlights

Effective strategy, strong demand

**Sold out** of all available capacity

Core AI infrastructure business

**19% adjusted EBITDA margin**

Group revenue growth

**+355% YoY  
+39% QoQ**

Core AI Infrastructure business grew at 400% YoY and 40% QoQ

Funding secured

**\$4.3 billion**

in convertible notes and an equity offering in September 2025

## Product update

### Nebius AI Cloud 3.0 “Aether”

Enterprise-grade security and compliance. Secured certifications like SOC 2, HIPAA, ISO 27001-series and more.

### Nebius Token Factory

Launch of production-scale inference platform bringing together inference, post-training and fine-grained access management with S3.

## Customer wins

Hyperscalers

 Microsoft

 Meta

Cutting-edge AI startups

 Black Forest Labs

 CURSOR

## Other businesses

Autonomous technologies

 AVRIDE

Secured strategic investment from Uber.

Over 215k deliveries YTD in US; on track for robotaxi launch in Dallas in December.

EdTech

 tripleten

2x revenue growth YoY supported by demand, efficient customer acquisition and product innovation.

Strategic equity stakes

 ClickHouse

Reported valuation of \$6.35bn in May 2025 funding round.

 Toloka

Data solutions business backed by Bezos Expeditions.

# Dear shareholders,

We are at the forefront of one of the most significant technological revolutions in history — and Nebius has quickly become a core enabler of the AI-driven economy. As you know, our mission is to become one of the world's leading AI infrastructure and cloud providers. Q3 2025 marked a major step forward in our trajectory, as we position the business for continued rapid growth in 2026 and beyond.

**During Q3, we made substantial progress on our key strategic objectives, while continuing to lay the foundations for future growth.**

- Once again, **we sold out of available capacity in Q3** — a consistent trend during the year — reflecting strong momentum and the effectiveness of our business strategy.
- In September, we also announced our **first large enterprise AI infrastructure win, with Microsoft**, valued between \$17.4 billion and \$19.4 billion. We are fully on track with the execution of this contract and expect revenue from this deal to start ramping up throughout 2026.
- Today we are pleased to announce our **second large AI infrastructure deal, this time with Meta**. Our agreement with Meta is valued at approximately \$3 billion, with a five-year term. Over the next three months, we plan to deploy the capacity needed to service the agreement. In fact, demand for this capacity was overwhelming, and the size of the contract was limited to the amount of capacity that we had available.
- Similar to our Microsoft deal, the economics of this agreement are attractive and will help accelerate the growth of our AI cloud business even further in 2026 and beyond.
- We will **continue to aggressively ramp up capacity** to be able to continue supporting these large deals, as well as to serve our core customers: startups and enterprises.
- On the financing front, **we raised a total of \$4.3 billion** in convertible notes and a follow-on equity offering in September to accelerate our investments in infrastructure, power, land, and GPUs.
- In addition to significant investments in our capacity footprint, we also continued investing in our AI cloud. To expand our addressable market opportunity among enterprise customers, **we released our new, enterprise-ready Nebius AI Cloud 3.0 "Aether" and our inference platform, Nebius Token Factory**. We believe the Aether release gives organizations the trust, control, and simplicity they need to run their most critical AI workloads. Nebius Token Factory is an inference platform that enables organizations to run open-source models in production with reliability, visibility, and control. We are continuing to build a large pipeline of new software and services, which we believe will provide deep moats and differentiate Nebius from other cloud companies.



Arkady Volozh,  
Founder and CEO

**"2025 has been a building year as we put in place the infrastructure and framework for future rapid growth. This year, we believe that we have successfully laid the foundations for an outstanding 2026 — a year that should firmly position us among the top AI cloud businesses globally. And at the same time, 2026 is still just the beginning."**

- Finally, we were pleased to welcome **Uber as a new strategic investor in Avride**. We believe this investment and partnership will help Avride to capitalize on the substantial market opportunity for autonomous driving technologies and services.

## Accelerating data center capacity

- Given the progress we've made to date, we are confident in our business strategy, operations, and overall market demand. The only real limitation on our revenue growth in 2025 has been the amount of capacity that we have been able to bring online. In the last few months, we have worked very hard to unlock this bottleneck, and we will continue doing so in 2026.
- Last quarter we guided for 1GW of contracted power by the end of 2026. We are currently in the process of securing additional sites that would bring our total contracted power (i.e., secured land with power) to approximately **2.5GW** by the end of 2026. Of this contracted amount, we expect to have **800MW to 1GW** of connected power (i.e., power connected to built data centers) by the end of 2026.

## Financing our growth

- To support our aggressive growth plans in 2026 and to maintain this pace of growth in 2027, we will utilize at least three sources of financing: corporate debt, asset-backed financing, and equity. We believe we will be able to secure asset-backed debt on attractive terms, supported by the creditworthiness of our largest customers. We will maintain a disciplined capital structure.
- On the equity front, we will be putting in place an at-the-market (ATM) equity program for up to 25 million Class A shares, and plan to file a prospectus supplement on November 12, 2025. We will evaluate the program regularly based on our capital needs. The program enables us to access equity funding on an efficient ongoing basis; however, we will remain dilution-sensitive as we prepare to finance future growth opportunities.

## Increasing revenue

- Given the scale of the financing that we expect to put in place in the next few months, and the extensive work that we have undertaken to build out our capacity, **we believe that we can achieve annualized run-rate revenue (ARR)<sup>1</sup> of \$7 billion to \$9 billion by the end of 2026.**

We have consistently said that we are committed to growing our business aggressively, and we are continuing to deliver on this commitment. 2025 has been a building year as we put in place the infrastructure and framework for future rapid growth. This year, we believe that we have successfully laid the foundations for an outstanding 2026 — a year that should firmly position us among the top AI cloud businesses globally. And at the same time, 2026 is still just the beginning.

Sincerely,

Arkady Volozh

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

# Group business update



## Driving scale

**We are on track for 220 MW of connected power and 100 MW of active power by the end of 2025**

We currently have data centers in three regions: North America, Europe, and the Middle East.

### United States

- We added capacity in Kansas City, and launched our first NVIDIA B200 cluster there.
- We announced that our new hyperscaler customer, Microsoft, will be served exclusively from our New Jersey data center.

### Europe

- The expansion of our Finland data center continued, and should be completed by the end of this year.
- In November we unveiled a deployment of NVIDIA B300s in the UK, supporting the UK's efforts to build out domestic compute capacity.

### Middle East

- In October we launched in Israel with a deployment of NVIDIA B200s.

**We now expect to have over 2.5GW of contracted power by the end of 2026, with 800MW to 1GW of connected power**

We are aggressively scaling our footprint to meet global demand for AI compute across the US, Europe, and Middle East.

- We plan to continue scaling our data centers in the UK and Israel, with the next phase of our UK expansion taking place in Q1 2026.
- We are also in the process of securing several large sites in the US and Europe; some of these have hundreds of megawatts of capacity.
- We also have the option to add hundreds of megawatts of incremental capacity in New Jersey in the future.

## Expanding and diversifying our customer base

**Our recent multi-billion dollar agreements highlight our ability to win the trust and confidence of some of the world's largest tech companies.**

- Partnerships like the ones we signed with Meta and Microsoft are an important part of our overall growth strategy. As well as driving scale, they also serve as powerful commercial and financial accelerators for our core AI cloud business.
- These wins with some of the world's most demanding AI leaders validate the strength of our offering and underscore how our investments in advanced AI infrastructure position Nebius as a preferred partner for major AI-centric enterprise customers.

**Nebius is trusted by the world's largest enterprises and cutting-edge startups alike.**

Our customers run their mission-critical AI workflows on Nebius because of our scale, robust technology, proven performance, and exceptional service. These differentiators are particularly important as we expand into the enterprise. We are seeing our investments in these capabilities start to have a meaningful impact on our growth trajectory.

- In the third quarter, we deepened relationships across industries — including scaled players like Shopify and World Labs — and onboarded innovative startups such as Cursor and Black Forest Labs.
- We are also winning new customers and expanding use cases in healthcare and life sciences, where partners such as Basecamp Research, Prima Mente, and Sword Health are building models to help solve some of the world's biggest medical challenges.
- We will continue to expand into additional sectors and use cases, investing to meet the advanced technical and compliance requirements of customers in physical AI, financial services, media and entertainment, and retail.

In this section, we discuss three forms of power: (1) contracted power, or power that is secured by contracts; (2) connected power, or power that is connected into data centers; and (3) active power, or power that is being consumed by IT equipment.

# Elevating cloud experience

Customers continue to choose Nebius for high-performance compute, native AI/ML orchestration, fast and scalable storage, and an expanding ecosystem of software integrations.

Our full-stack AI architecture provides the flexibility, control and efficiency that large-scale enterprises and research organizations demand, while also delivering the speed required by nimble AI startups.

Nebius is building the infrastructure needed to power the AI market — whether for training and refining large language models, or to enable fast and efficient inference.

## Delivering enterprise-readiness with Nebius AI Cloud 3.0 "Aether"

Nebius AI Cloud 3.0 "Aether" is the latest major release of Nebius AI Cloud that addresses the unique needs of enterprises running business-critical workloads at scale in highly regulated environments. With version 3.0, we have added:

### Enterprise-grade security & compliance

- Received SOC 2 Type II including HIPAA, and ISO 27001 certifications validated by independent auditors, enabling AI deployment in highly regulated contexts including healthcare, financial services, and government.
- Aligned with NIS2, DORA, ISO 27032, ISO 27701 and ISO 27799 regulatory frameworks for security.
- Embedded secrets manager for secure API key and credential management.

### Advanced governance & control

- Enhanced Identity and Access Management (IAM) features, including more granular controls for policy enforcement and collaboration across different teams and users.
- Created a comprehensive observability suite with searchable logs, histograms, metrics ingestion, and auditing, bringing together metrics from different services under one consolidated view.

### Developer productivity

- Updated user interface (UI) and navigation that better align with AI/ML developer workflows, including a simplified application catalogue.
- Developed support for VM creation with preloaded applications for rapid AI workload deployment.
- Simplified GPU, CPU, and memory quota allocation and billing.
- Expanded ecosystem support including SkyPilot Server and Anyscale integrations.

**We continued to make significant progress in improving our AI cloud performance through advances in speed, reliability, and flexibility.**

We believe our platform continually demonstrates best-in-class performance, as validated by key third party benchmarks.

- During the quarter, we achieved a new inference performance milestone. Nebius achieved leading performance results in the latest submission of MLPerf Inference v5.1 benchmarks for three AI systems accelerated by the most in-demand NVIDIA systems on the market.
- Nebius AI Cloud set a new peak in performance delivered by NVIDIA GB200 NVL72 systems. Compared to the best results of the previous round, our one-host installation achieved a ~7% and ~14% performance increase running Llama 3.1 405B inference in offline and server modes, respectively.
- In another submission running Llama 3.1 405B inference on NVIDIA HGX B200, Nebius AI Cloud recorded a 3x and 4.3x increase in tokens per second compared to our NVIDIA H200 results in this MLPerf submission in offline and server scenarios, respectively.
- We also delivered significant increases in storage performance: file storage write speeds doubled, read speeds increased by 50%, and WEKA-based storage now reaches approximately 20 GBps read / 18 GBps write speeds.
- In terms of reliability and performance, we added additional active health checks, as well as self-healing node capabilities in our Managed Kubernetes and Managed Soperator components. We believe these improvements make our platform one of the most reliable and performant in the market.

## Delivering the highest bare-metal performance standards on virtualized cloud infrastructure

Nebius is one of the first vendors to achieve NVIDIA Exemplar Status for training workloads running on NVIDIA H200 GPUs. This recognition validates Nebius's performance against some of the industry's most rigorous standards for performance, resiliency, and scalability, within a virtualized, flexible cloud environment.

## Launching Nebius Token Factory to serve enterprise needs

- In November, we launched Nebius Token Factory, which unifies large-scale inference, model optimization and fine-tuning in one governed platform. Token Factory helps vertical AI product builders, ISVs and digital enterprises deploy open-source models at scale with predictable latency, transparent cost, and built-in compliance.
- At its core, Token Factory is a production-scale inference platform, which serves and optimizes open-source AI models at scale and brings together high-performance inference, post-training, and fine-grained access management so organizations can run those models in production with reliability, visibility, and control.

## Enterprise-ready

- Enterprises can deploy and scale models such as OpenAI OSS, Qwen, DeepSeek and many others on dedicated endpoints with guaranteed performance and isolation, tuned for sub-second latency and delivering 99.9% uptime.
- Nebius Token Factory gives enterprises a single platform to industrialize the post-training lifecycle: transforming open-source checkpoints into optimized, production-ready systems with guaranteed performance and transparent cost per token.

## Optimization without overhead

- Token Factory integrates fine-tuning and distillation pipelines that let teams adapt large open-source models to proprietary data, while significantly cutting inference costs, with one major client reporting a 26-fold reduction.
- Optimized models can be deployed directly to production endpoints with no manual infrastructure work required, giving enterprises faster iteration and lower cost per token.

## Supported by in-house AI R&D capabilities

- We're on a mission to help builders get production-grade results from open-source models, and we're investing heavily in the research that drives our core products.
- Our SWE-rebench dataset has achieved remarkable adoption, becoming the most downloaded dataset on HuggingFace. Leading AI labs such as Moonshot (the developer of Kimi K2) are now leveraging it as their primary source for training reinforcement learning models in software engineering environments.
- We launched and maintain a continuously updated leaderboard at [swe-rebench.com](http://swe-rebench.com), which has gained substantial traction within the research community and is updated monthly, ensuring that the task sets and model benchmarks remain current and relevant.
- Beyond benchmarking, we have contributed to the development of AI agents, including a recent paper that introduces a novel methodology for applying reinforcement learning to long-context, multi-turn AI agents, addressing a critical challenge in building more capable autonomous systems.

## Expanding our ecosystem

**Nebius continued to expand its global partner ecosystem this quarter, advancing both market reach and enterprise adoption. Strategic alliances and platform integrations strengthened Nebius's position as one of the key infrastructure layers of choice for AI innovation.**

### Channel expansion

- In October, TD Synnex announced a partnership with Nebius as its preferred AI Infrastructure-as-a-Service partner in North America, to bring Nebius's high-performance AI cloud to the IT ecosystem.

- This partnership provides access to thousands of resellers and enterprise buyers, creating a scalable channel to deliver Nebius infrastructure and services across new markets and verticals.

### Platform integrations and ecosystem growth

Nebius also deepened its integration ecosystem with leading AI development platforms to provide end-to-end solutions from procurement to deployment.

- **Baseten:** Launched AI video and inference stack on Nebius, delivering predictable P90/P99 latency, high utilization, and stable uptime for production-grade workloads.
- **Anyscale & SkyPilot:** Allow Nebius customers to seamlessly scale and orchestrate distributed AI training and inference on Nebius's GPU infrastructure, reducing operational complexity while ensuring predictable, high-performance workloads.
- **Qdrant & dstack:** Expanded access through RAG content tools and GPU marketplaces.
- **Union.ai:** Joined as a new platform partner to support AI and data workflow orchestration at scale.

These partnerships collectively expand Nebius's reach, strengthen brand awareness, and accelerate time-to-value for customers, reinforcing our position as the preferred infrastructure foundation for scalable, sovereign, enterprise-grade AI.



## Avride

**In October, Avride announced a strategic partnership with Uber, under which Uber and Nebius will make strategic investment and other commitments of up to \$375 million in Avride.**

- The investment enables significant and rapid scaling of Avride's business to capitalize on the substantial market opportunity for autonomous driving technologies and services.
- The deal is validation of Avride's strong technical capabilities in developing autonomous driving and navigation technologies.
- The investment builds on Avride's commercial partnership with Uber, following the signing in 2024 of a multi-year strategic agreement.
- Nebius Group continues to support Avride's strategic direction and long-term value creation, and intends to retain a significant stake in the company, securing substantial potential upside for Nebius shareholders.

## Preparing for ride-hailing launch in Dallas

- In Q3, Avride finalized pre-launch testing of its autonomous vehicles in Dallas and integration with the Uber ride-hailing platform.
- Dallas is set to become the first city to host Avride and Uber's public robotaxi service in December 2025.

## Robodelivery continues to scale at pace after achieving important milestone

- Avride's U.S. delivery business continues to grow, surpassing 215,000 deliveries YTD.
- In October, with expanded delivery zones in Jersey City, Dallas, and Austin, along with growth in campus operations, daily deliveries exceeded 1,900. Additional cities are expected to launch on the Uber Eats platform in the coming months.
- Avride began working with major restaurant chains — Shake Shack, Wendy's, and White Castle — and continues to expand its list of restaurant partners.
- The business recently expanded to the University of Arizona in Tucson with a fleet of 60+ delivery robots. The company expects to scale its successful campus operating model to additional universities in the coming months.
- In October 2025, Avride announced the launch of robot delivery on the Grubhub marketplace — expanding the partnership beyond campuses for the first time. This marks Grubhub's first use of autonomous delivery in its main marketplace and Wonder's first robot delivery service for its customers.
- In Japan, Avride's partnership with Mitsui Fudosan to provide autonomous retail logistics deliveries at the country's largest outlet mall has been successful and continues to expand.



## TripleTen

### New student enrollment driving revenue

- TripleTen continued to deliver strong revenue growth of approximately 100% year over year during the quarter, supported by student enrollment which increased by approximately 6,000 new students and growth in average fees per student.
- These trends reflect sustained demand across core markets, including the US and Latin America. Customer acquisition remained efficient, supported by continued referral-driven enrollments and sales team productivity.

## New products and services are gaining traction

- Launched a new program called AI Analyst to help students retool their skills for AI.
- Launched a new B2B product, Nebius Academy, where TripleTen will sell into organizations to provide courses for their employees.

## Continuing to utilize AI to drive further operational efficiency

- TripleTen is now leveraging AI for various student services, like resume building, coaching, and more.
- Furthermore, TripleTen is leveraging AI to help build our course material and videos, which provide significant cost savings when compared to traditional methods of content creation.

## Equity investments: Toloka and ClickHouse

In addition to our non-core businesses, we own equity stakes in both Toloka and ClickHouse.

### Toloka

Toloka is a leading data provider for LLM and GenAI developers, delivering scalable, high-quality data solutions for all stages of AI development.

Although we no longer hold voting control of Toloka, we maintain a significant majority economic interest in the business and are encouraged by the growing investor interest in the AI data provider space, reinforcing our confidence in the substantial growth opportunities ahead.

In Q3, Toloka completed the core technology for its hybrid agent, which combines AI with human expertise to execute long-horizon tasks, RL Gyms and large-scale data production. This advancement enables Toloka to serve a broader customer base, including leading AI labs, global enterprises, researchers, and individual users.

Toloka plans to launch Tendem.ai, a unique hybrid human-AI agent, in the near future. Tendem.ai integrates advanced automation with expert human oversight to deliver accurate, reliable, and ready-to-use results across a wide range of applications. For a given set of tasks, Tendem.ai agent can accomplish a large portion of those tasks, enabling a person to finish the last mile of those tasks.

### ClickHouse

We currently own a minority equity stake in ClickHouse, a leader in real-time analytics, data warehousing, observability, and AI/ML, that we spun out in 2021.

# Financial update

Comparisons refer to the year-over-year changes from Q3'24 to Q3'25, unless otherwise indicated.

## Revenue

In USD \$ millions

Three months ending September 30

	2024	2025	Change
Revenue	32.1	146.1	355%

Group revenue was \$146.1 million in Q3, up 355% and up 39% quarter-over-quarter, excluding Toloka's results, which have been reclassified to discontinued operations.

- Our core AI infrastructure business accounted for approximately 90% of total group revenue in Q3'25. Annualized run-rate revenue (ARR)<sup>1</sup> at the end of September was approximately \$551 million.
- Demand remained strong, driven by increasing AI use cases, growth in adoption, and the expansion of our client base.
- The expansion of our sales force and customer support teams continued to propel strong growth. Even with our significant GPU capacity additions, we operated at peak utilization.

- Other businesses accounted for the remainder of the Group's revenue.
  - TripleTen continued to deliver strong revenue growth of 100% supported by the addition of approximately 6,000 new students across the US and Latin America and the launch of new programs.
  - Avride's revenue contribution in Q3'25 was immaterial.

Note: Comparisons refer to the year-over-year changes from Q3'24 to Q3'25, unless otherwise indicated.  
Toloka results have been reclassified to discontinued operations in the current and prior-year periods.

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

## Operating expense

In USD \$ millions

Three months ending September 30

Expense category	2024	2025	Change
Cost of revenues	9.9	42.9	333%
as a percentage of revenues	31%	29%	
Product development	31.4	44.9	43%
as a percentage of revenues	98%	31%	
Sales, general and administrative	47.9	89.5	87%
as a percentage of revenues	149%	61%	
Depreciation and amortization	23.5	99.0	321%
as a percentage of revenues	73%	68%	
<b>Total operating costs and expenses</b>	<b>112.7</b>	<b>276.3</b>	<b>145%</b>
as a percentage of revenues	351%	189%	

**Cost of revenue** was \$42.9 million, up 333%, representing 29% of revenue, down from 31% in Q3 2024. The dollar increase was due to the expansion of our core AI infrastructure business, mainly driven by costs incurred for co-location and operating lease agreements, as well as the hiring of additional personnel attributable primarily to data-center operations. Our Avride and TripleTen businesses had a limited contribution to the overall increase in cost of revenues.

**Product development expenses** were \$44.9 million, up 43%, representing 31% of revenue, a significant decline from the 98% of revenue reported in Q3'24. The dollar increase in product development expenses was primarily driven by the expansion of our product development teams.

**Sales, general and administrative expenses** were \$89.5 million, representing 61% of revenue, which was down from 149% of revenue in Q3'24. In Q3'24, we incurred significant consulting, legal

and professional fees which were incremental to our main operating activities. Excluding these incremental expenses, SG&A dollar expenses in Q3'25 increased from Q3'24 due to an increase in share-based compensation expense allocated to personnel engaged in sales, general and administrative activities, followed by the growth in marketing and advertising expenses.

**Depreciation and amortization expenses** ("D&A") for the group were \$99.0 million compared to \$23.5 million in the prior year period, or 68% of revenue for the period, down from 73% in Q3'24. The key driver of the dollar increase in D&A expenses was the significant ramp-up in investments in GPU-related capital expenditures and related data center hardware for the core AI infrastructure business. We depreciate our hardware over a four-year period, which we believe is conservative when compared to common industry practices.

Note: Comparisons refer to the year-over-year changes from Q3'24 to Q3'25, unless otherwise indicated. Toloka results have been reclassified to discontinued operations in the current and prior-year periods.

## Adjusted EBITDA

In USD \$ millions

Three months ending September 30

	2024	2025	Change
Adj. EBITDA (loss)	(45.9)	(5.2)	-89%
as a percentage of revenues	-143%	-4%	

The Group's adjusted EBITDA loss narrowed significantly to \$5.2 million in Q3'25, representing a \$40.7 million year-to-year improvement. This was driven primarily by the growth of our core

business, which generated positive adjusted EBITDA in Q3'25. Our investments in Avride and TripleTen resulted in the reported adjusted EBITDA loss for the quarter.

## Capital expenditures

In Q3 2025, our capital expenditures were \$955.5 million, primarily driven by purchases of GPUs and GPU-related hardware, and our data center expansion activities. In addition to reporting capital expenditures as we have in the past, this quarter we began providing a full Statement of Cash Flows.

## Guidance

The company will share detailed year-end 2025 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 November 11, 2025 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](https://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, 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. The potential risks and uncertainties that could cause actual results to differ from the results predicted or implied by such statements include, among others, our ability to successfully compete in our sector; to implement our business plans; to continue to successfully capture customers; to continue to successfully obtain required supplies of hardware on acceptable terms; and to obtain further debt or equity financing that may be necessary to achieve our objectives on acceptable terms. Many of these risks and uncertainties depend on the actions

of third parties and are largely outside of our control. We also continue to be subject to many of the risks and uncertainties included 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, 2024 filed with the U.S. Securities and Exchange Commission ("SEC") on April 30, 2025, which are available on our investor relations website at <https://group.nebius.com> and on the SEC website at [www.sec.gov](http://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.

## 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 plus (1) depreciation and amortization, (2) certain SBC expense, (3) interest expense, (4) income tax expense/(benefit), (5) one-off restructuring and other expenses, less (1) interest income, (2) other income/(loss), net, (3) income/(loss) from equity method investments and (4) gain from revaluation of investment in equity securities.
- Adjusted net income/(loss) means U.S. GAAP net income/(loss) from continuing operations plus (1) certain SBC expense, (2) one-off restructuring and other expenses, (3) amortization of debt discount and issuance costs, net of interest expense capitalized less (1) foreign exchange gains and (2) gain from revaluation of investment 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 certain 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, 2024 filed with the U.S. Securities and Exchange Commission ("SEC") on April 30, 2025. In Q2 2025, following the completion of the investment transaction in Toloka, an AI development platform, Nebius ceased to hold majority voting power in Toloka and no longer includes Toloka's results in Nebius' consolidated financial statements and reports its stake as equity method investment. The Toloka's results for prior periods were reclassified to discontinued operations.

### Certain 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. However, because we settled some RSU equity awards of our employees granted before 2022 in cash during 2024, a portion of stock-based compensation expense for 2024 was included in Adjusted EBITDA/(loss).

### Foreign exchange gains/(losses)

The functional currency of Nebius Group N.V. is the United States Dollar, which is also the Group's current 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 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, plant 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.

The tables at the end of this release provide detailed reconciliations of each non-GAAP financial measure we use from the most directly comparable U.S. GAAP financial measure.

Contact investor relations:  
[askIR@nebius.com](mailto:askIR@nebius.com)

**NEBIUS**