

Nebius Group Letter to shareholders Q2 2025

Highlights Q2 2025

Financial performance

Forecasted ARR¹ for core business increased

**\$900 million –
\$1.1 billion**

Owing to strong momentum, up from
\$750 million – \$1 billion announced previously

Core AI infrastructure business

**Positive
adjusted EBITDA**

Group adj. EBITDA loss \$21.0M;
\$37.1M YOY improvement

Group revenue

\$105.1 million

+106% QoQ; +625% YoY

Cash at the end of Q2

\$1.68 billion

CAPEX

\$510.6 million

driving capacity expansion

Capacity expansion

220 MW

Connected power²

CYE 2025

>1 GW

Contracted power

CYE 2026

Delivery timeline

Q2'25

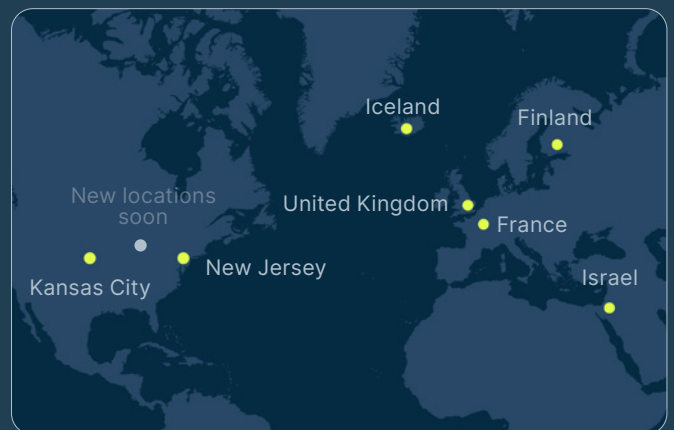
Kansas City
UK & Israel
announced

CYE 2025

NJ Phases I and II
UK & Israel
Finland expansion

CYE 2026

NJ Phase III
New greenfield sites
Finland expansion
completed



Product and ecosystem update



Increased speed.
Top-tier speed and **scalability** of clusters as verified
by MLPerf® Training v5.0
Benchmark Results



Enhanced reliability.
More uptime and **lower costs** for AI workloads
with autohealing and health
checks



Storage and flexibility.
Storage optimization for-
streaming & checkpoint
ops, **compatibility with S3**

New partners
in the ecosystem

Anyscale

baseten

Lightning ^{AI}

MITHRIL

SkyPilot

Selected enterprise customers

CLOUDFLARE

prosus

shopify

Other businesses

AURIDE

Among partners:

Uber GRUBHUB

HYUNDAI

tripleten

AI-ready, AI-powered

+6,000

New students in Q2

Equity investments

ClickHouse

Reported valuation
of \$6bn in latest
funding round

Toloka

Data solutions business
backed by Bezon
Expeditions

(1) Annualized run-rate revenue — revenue for the last month of the period multiplied by 12

(2) Includes both active power and power that is fully provisioned and can be activated immediately upon GPU installation

Fellow shareholders,

These are exciting times. With AI developing at pace, we stand at the brink of the next great leap in technology. To achieve this, we need to build entirely new infrastructure for AI. This is a great challenge for the next decade and beyond.

Outside of a handful of big tech companies, only a few newcomers have both the technological expertise and the multi-billion-dollar scale to play in this field. Nebius is one of them.

Demand is continuing to grow rapidly as frontier AI labs build more LLMs, thousands of new AI-native startups develop ever greater numbers of applications, and large enterprises incorporate AI into critical workstreams. We expect the fundamental trends in our space to continue to drive growth for years to come.

Nebius is a recognized leader in AI infrastructure — across compute, software and services — and we are positioning the company to win.

In our most recent quarter, we had a number of important accomplishments:

We continued to deliver exceptional results across the board.

- We grew revenue 625% year-over-year and 106% quarter-over-quarter – a significant acceleration from Q1 – with growth supported by strong customer demand and near-peak utilization of our platform.
- We achieved positive adjusted EBITDA in our core business sooner than expected.

As a result of this strong momentum, we are raising our annualized run-rate revenue outlook for the year end from \$750 million-\$1.0 billion to \$900 million-\$1.1 billion.



Arkady Volozh,
Founder and CEO

“Demand for AI infrastructure — compute, software and services — is only going to get stronger as use cases multiply. We are aggressively scaling up capacity to capture this substantial opportunity, and are in the process of securing more than 1 GW of capacity by the end of 2026.”

We are aggressively scaling capacity and targeting 1 GW of power in 2026.

- By the end of 2025, we expect our connected power to reach 220 MW. This includes capacity that is or will be connected to our data centers.
- In Finland, we are currently expanding up to 75 MW. In New Jersey, we recently increased our commitment to secure 100 MW of power this year and 100 MW of power in 2026.
- Beyond that, we are in advanced discussions for two significant new greenfield sites in the US. Overall, we are in the process of securing more than 1 GW of power by the end of 2026.

We continued to expand and diversify our customer base.

- We are gaining traction on the enterprise side, adding large global technology logos including Cloudflare, Prosus and Shopify.
- We also remained one of the top neocloud providers for AI-native startups, winning customers like HeyGen, Lightning.AI, Potoroom and many others.

We made significant enhancements to our cloud platform and customer experience.

These enhancements are designed to provide customers with greater performance, improved reliability and better flexibility, including:

- Achieving faster run-time speed of our clusters as verified by MLPerf® Training v5.0 submission
- Improving cluster reliability by significantly increasing Mean Time Between Failure
- Launching Managed Soperator, our fully managed Slurm solution, available as a one-click, self-service option
- Strengthening our storage offering with a new class-of-object storage and third-party solutions
- Extending our integrations with leading AI development platforms and services
- We also continue to invest in new products like our Inference-as-a-Service solution to help customers utilize the latest LLM models for their inference workloads.

We continued to grow our ecosystem.

- Our strategy is to build an API-first, developer friendly Platform-as-a-Service.
- We added new integrations and built on existing ones that make Nebius infrastructure more accessible to AI developers and researchers, including with AnyScale, BaseTen, Mithril (Foundry), LightningAI, and we deepened our collaborations with SkyPoint and NVIDIA.

We are in the midst of a once-in-a-generation opportunity.

Demand for AI infrastructure — compute, software and services — is only going to get stronger.

Looking forward, we are fortunate to have multiple levers to finance our aggressive growth plans. To date, we have secured more than \$4 billion in capital. Our strong balance sheet provides us with a wide range of potential financing options to support our ambitions. Furthermore, we believe that our portfolio of businesses and equity investments should provide us with additional billions of dollars in incremental capital over time, enabling us to continue investing in our core business.

I'd like to thank our employees for their contributions to our current success, and to our shareholders for their continued support.

Sincerely,
Arkady Volozh

Group business update



NEBIUS

Driving scale

We plan to have 220 MW of power connected to our data centers by the end of this year.

At the end of Q2, our footprint spanned five active locations globally. In addition, we announced plans to launch in the UK and Israel later this year to support local AI development strategies and create new opportunities for British and Israeli businesses, researchers and public sectors. By the end of this year, we will operate seven AI clusters in six countries across Europe, the US and the Middle East, making us one of the largest independent AI infrastructure builders globally. We expect to have 220 MW of connected power, which includes both active power and power that is fully provisioned in our data centers and can be activated immediately upon GPU installation.

United States

Strong demand continues to outpace capacity in the US as AI use-cases multiply. In light of the current market demand, coupled with capacity constraints, we are continuing to significantly grow our US capacity in 2025.

Our Kansas City data center officially launched in April and we added more compute capacity in early July.

Our new facility in Vinefield, New Jersey is on track

to open by the end of the summer and we plan to offer instances from this facility this fall.

We now have 200 MW of contracted power that is in active construction in New Jersey, of which 100 MW will be made available this year, with the rest following in 2026.

Europe & Middle East

Europe and the Middle East are large and thriving AI markets, with significant venture capital funding for AI startups and a growing number of scaled enterprises seeking to leverage AI to power their products and processes. Nebius continued to expand its global capacity footprint during the quarter, cementing our strong positions in both regions.

We announced the development of two new data centers in the UK and Israel to support the large growth opportunities in Europe and the Middle East.

In June, we announced our first general availability of NVIDIA GB200 Grace Blackwell Superchip capacity for customers in Europe. We expect our deployment of NVIDIA Blackwell Ultra GPUs in the UK to be operational during Q4 2025.

We are on track to substantially complete the expansion of our Finnish data center by the end of 2025, bringing approximately 75 MW of contracted power to that facility.

Expanding & diversifying our customer base

Our full-stack AI Cloud offering, brand and partnership efforts, coupled with our sales execution, are starting to deliver strong results.

Q2 was an important turning point for our business. We continued to make strong progress with our core base of AI native tech startups, and also started to add new enterprise logos like Cloudflare, Prosus and Shopify. Customers are taking note of our differentiated offering, our 24/7 customer and technical support, and our ability to build products and tools tailored to our customers' needs.

In Q2, we remained a top neocloud provider for native AI tech startups, such as HeyGen, Lightning AI, Photoroom, and many others.

We also continue to serve customers across a broad range of industries, including financial services, life sciences, media & entertainment and technology.

Notably, Shopify recently discussed at length how they are using AI in their business, and how Nebius helps power their AI infrastructure needs. Read their insightful blog [here](#).

Elevating cloud experience

Our AI Cloud is a growing and differentiated platform.

From high-performance compute and orchestration (Slurm, Kubernetes) to ML operations, scalable storage, and a growing ecosystem of third-party AI tools, Nebius continues to lead as a full-stack AI cloud platform.

Our AI-native architecture gives customers the flexibility of Infrastructure-as-a-Service (IaaS) with the velocity of Platform-as-a-Service (PaaS), and our momentum this quarter reflects the growing demand for performant, production-grade infrastructure that scales with their ambition.

The vast majority of our contracted customers use our AI Cloud because we provide mission-critical software and services that provide them with the applications and tools they need to manage and grow their AI workloads.

This quarter, we launched Nebius AI Cloud v3, which delivered a set of infrastructure and software

upgrades focused on performance optimization, scheduling efficiency, system reliability, and customer flexibility at scale.

Setting new standards in performance

We continue to be a leader in infrastructure performance. In our latest MLPerf® Training v5.0 submission, we achieved top-tier speed and scalability. Our benchmarks showed near-linear scaling — a ~2x speedup when doubling from 64 to 128 nodes — a reflection of the powerful performance and efficiency of the Nebius infrastructure.

Complementing this, we shipped topology-aware job scheduling across multiple frameworks (Slurm, Kubernetes, Volcano, Kueue). We also integrated NVIDIA Topograph to streamline topology data ingestion.

Additionally, our ISEG2 system ranked #13 in the June 2025 TOP500 list of the world's fastest supercomputers, and #1 in Europe among commercially available systems.

Reliability at scale

We pushed our infrastructure reliability to new heights, significantly increasing Mean Time Before Failure (MTBF) to 168,000 GPU hours per week for a 3,000-GPU cluster. This is the result of multiple improvements:

- We reduced recurring problems that lead to job failures by 60%. We accomplished this by implementing a software stack that aggregates detailed analytics across all clusters automatically, enabling us to identify and focus on the most impactful root causes more effectively.
- Smarter health checks and improved autohealing for Slurm clusters and Kubernetes environments, allowing for seamless node recovery and uninterrupted training. This means significantly more uptime and lower costs for AI workloads, allowing customers to complete their runs in a timely manner for less.

Expanding flexibility and customer choice

To simplify data onboarding and broaden support for diverse architectures, we launched:

A new data migration service for easy ingestion from any S3-compatible source – no extra infrastructure required. We believe this will make it much easier for customers to work with Nebius going forward.

Support for third-party AI-optimized storage solutions like WEKA and VAST, giving customers more control and performance tuning capabilities.

A new class of Object Storage optimized for streaming and checkpointing workloads, delivering up to 2 GB/s write throughput per GPU.

Accelerating our inference platform

Over the next several years, we believe inference has the potential to become a bigger opportunity than training, as more companies bring AI solutions at scale into their businesses.

We support AI-centric integrated software vendors and enterprises that are applying AI in critical workflows with our fully vertically integrated inference-as-a-Service product. To do this, we are building an enterprise-grade platform that enables the deployment and scaling of open-source AI models like Llama, Qwen, DeepSeek, Flux, OpenAI's new models, and others, with high performance and reliability on our dedicated infrastructure, giving full control in a unified API.

Our platform runs on top of Nebius' proven, scaled infrastructure to solve the biggest pain points in production AI: unpredictable latency, GPU bottlenecks, and expensive, inflexible platforms.



Roman Chernin, our Chief Business Officer, will now focus on new business innovation, starting with scaling our Inference-as-a-Service platform. We have a shared commitment to scaling inference into a key pillar of the business and to capture its significant market potential.

Expanding our ecosystem

We continued to expand our reach and drive brand affinity and awareness through strategic partnerships that strengthen our platform, increase visibility, and broaden our global presence.

We deepened our collaboration with NVIDIA by expanding our integrations with NVIDIA's accelerated computing platform, including support for the NVIDIA AI Enterprise software stack. In parallel, we added new integrations that make Nebius infrastructure more accessible to developers and researchers:

AnyScale: offers a managed platform, built on Ray, an open-source unified compute framework for scaling AI and Python applications. This platform simplifies the deployment, scaling, and management of Ray applications. Our full integration with Anyscale facilitates easier deployment through the Nebius marketplace.

BaseTen: operates a high-performance inference platform serving some of the leading AI companies in the market and early adopting enterprises. We completed full integration with the BaseTen inference platform in Q2. BaseTen now runs internal workloads on Nebius, and end users can deploy on BaseTen Cloud with Nebius as the underlying infrastructure

LightningAI: provides an end-to-end platform designed to simplify and accelerate the development, training, and deployment of AI models and application development. Lightning AI integrated with Nebius, enabling its 300,000+ community members to access our on-demand and preemptible compute offerings.

Mithril (Foundry): is a platform designed to simplify and optimize the process of building, training, and deploying machine learning models. Nebius infrastructure is now available through Mithril's omnicloud platform, giving users access to Nebius compute capacity in both US and EU regions for training, inference, and experimentation.

SkyPilot: offers an open-source framework for running AI and batch workloads on any infrastructure. We expanded our partnership with SkyPilot from Q1, and our April integration gives users a seamless path to run workloads on Nebius AI Cloud, helping bring new users to the platform while offering existing SkyPilot users an easy way to leverage Nebius's



Avride

Expanded robotic delivery use cases across multiple geographies.

Avride continued to scale its autonomous delivery operations. In the US, robots delivered Uber Eats orders from restaurants across Jersey City, Dallas and Austin. The business also piloted a grocery delivery use case from an H-E-B supermarket, one of Texas's largest grocery chains. Restaurant delivery operations also continued at The Ohio State University through a partnership with Grubhub, launched in early 2025.

In Japan, Avride partnered with Mitsui Fudosan to deploy robots at the country's largest outlet mall. Robots deployed as part of this pilot are tasked with handling warehouse-to-store logistics, such as delivering packing materials to stores and routing packed items to central pickup points for delivery.

Advanced preparations for autonomous ride-hailing launch in Dallas.

In Q2, Avride continued fleet testing in Dallas in preparation for its commercial launch with Uber later in 2025. Testing focuses on real-world operating conditions to validate safety and performance. The fleet includes Ioniq autonomous vehicles developed in partnership with Hyundai, with Avride responsible for software development and systems integration.



TripleTen

Strong growth in student enrollment driving revenue.

TripleTen continued to deliver strong growth, supported by a quarter-over-quarter enrollment increase of approximately 6,000 new learners and growth in average revenue per check. 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.

Driving learning outcomes with a focus on margins.

TripleTen further leveraged AI-based tools to improve learning outcomes while optimizing operational efficiency. The integration of AI Tutor contributed to a nearly 35% increase in project completion rates, while reducing costs associated with reliance on manual instruction. To support strong placement outcomes, TripleTen expanded job support services with no impact to cost. This included introducing career preparation earlier in the learner journey, doubling externship participation, and adding CV placement specialists to improve interview conversion. In Q2, TripleTen also launched a new AI Automation course as part of its broader effort to stay ahead of emerging workforce demands.

ClickHouse and Toloka

In addition to our smaller business units, we own equity stakes in both Toloka and ClickHouse.

Toloka

In May, Toloka raised capital from leading investors Bezos Expeditions and Mikhail Parakhin, as well as Nebius. In connection with this transaction, Nebius Group's voting share in Toloka dropped below 50%, giving the Toloka team more autonomy to grow and scale their business. Nebius continues to own a significant majority economic stake in the business and intends to continue to support its growth and success. Our remaining stake in Toloka is accounted for under the equity method.

ClickHouse

ClickHouse is an open-source database platform, which we originally spun out in 2021. We currently own a minority equity stake in ClickHouse, and participated in the company's most recent fundraising round.

Committed to scaling sustainably

Nebius published its 2024 sustainability report, underscoring how sustainability delivers wins for the business, customers and society more broadly.

The report – Nebius's first prepared with reference to the European Sustainability Reporting Standards – outlines how sustainability is embedded in our core strategy. It demonstrates how this focus translates directly into competitive advantages, including operational cost leadership through energy efficiency, readiness for evolving regulatory requirements, and differentiation in an increasingly sustainability-conscious market.

In the reporting year, we achieved approximately 20% lower total cost of ownership through infrastructure efficiencies, including resilient hardware design and advanced energy-saving solutions. Our cus-

tom-designed servers alone saved 10 GWh of energy compared to off-the-shelf alternatives. Additionally, our data center in Finland covered 65% of the local municipality's heating needs through heat recovery, while operating cooling systems without water or refrigerants.

In total, 94% of the electricity we consumed came from low-carbon sources, resulting in an emissions intensity of just 0.04 tCO₂-eq/MWh – among the lowest in the tech sector and reflecting our commitment to sustainable, climate-responsible operations.

For more information, please see our 2024 Sustainability Report at nebius.com/sustainability.

Financial update

Revenue

In USD \$ millions	Three months ending June 30		
	2024	2025	Change
Revenue	14.5	105.1	625%

Group revenue was \$105.1 million, up 625% and up 106% quarter-over-quarter excluding Toloka’s results, which have been reclassified to discontinued operations.

Our core AI infrastructure business, which accounted for the vast majority of total group revenue in Q2’25, ended the quarter with annualized run-rate revenue of \$430 million, up from \$249 million at the end of March.


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 near-peak utilization by the end of the quarter.

Other businesses accounted for the remainder of the Group’s revenue.

- TripleTen continued to deliver strong revenue growth 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 Q2 was immaterial.

Group* revenue, \$ million



Note: Comparisons refer to the year-over-year changes from Q2’24 to Q2’25, unless otherwise indicated. Toloka results have been reclassified to discontinued operations in the current and prior-year periods.

(*) Nebius, TripleTen and Avride

Operating expense

In USD \$ millions

Three months ending June 30

Expense category	2024	2025	Change
Cost of revenues	7.7	30.1	291%
as a percentage of revenues	53%	29%	
Product development	32.0	42.8	34%
as a percentage of revenues	221%	41%	
Sales, general and administrative	75.6	68.2	-10%
as a percentage of revenues	521%	65%	
Depreciation and amortization	11.4	75.2	n/m
as a percentage of revenues	79%	72%	
Total operating costs and expenses	126.7	216.3	71%
as a percentage of revenues	874%	206%	

Cost of revenue was \$30.1 million, up 291%, representing 29% of revenue, down from 53% in Q2 2024. The decline as percentage of revenue was primarily driven by the scaling of our core AI infrastructure business.

Product development expenses were \$42.8 million, up 34%, representing 41% of revenue, a significant decline from the 221% of revenue reported in Q2'24. The dollar increase in product development expenses was primarily driven by the expansion of our engineering teams.

Sales, general and administrative expenses were \$68.2 million, down 10%, representing 65% of revenue, which was down from 521% of revenue in Q2'24. In Q2'24 we incurred significant consulting, legal

and professional fees which were incremental to our main operating activities. Excluding these incremental expenses, SG&A in Q2'25 increased from Q2'24 due to higher share-based compensation expenses.

Depreciation and amortization expenses ("D&A") for the group were \$75.2 million compared to \$11.4 million in the prior year period, or 72% of revenue for the period, down from 79% in Q2'24. 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 some other cloud providers.

Adjusted EBITDA

In USD \$ millions

Three months ending June 30

	2024	2025	Change
Adj. EBITDA (loss)	(58.1)	(21.0)	-64%
as a percentage of revenues	-401%	-20%	

The Group's adjusted EBITDA loss was \$21.0 million in Q2'25, representing a \$37.1 million year-to-year improvement. This was driven primarily by the growth of our core business, which generated posi-

tive adjusted EBITDA in Q2'25. Our investments in Avride and TripleTen resulted in an adjusted EBITDA loss for the quarter.

Capital expenditures

In Q2, our capital expenditures were \$510.6 million, primarily driven by purchases of GPUs and GPU-related hardware, and our data center expansion activities.

Guidance

The company will share detailed guidance on its Q2 2025 earnings call.

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 7, 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.

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 needs 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 guarantees of future performance. Actual results may differ materially due to various risks and uncertainties, including, but not limited to, our ability to successfully compete in our sector; implement our business plans; continue to successfully attract and retain customers; contin-

ue to successfully secure necessary hardware and supplies; and to obtain additional financing, that may be necessary to achieve our objectives, on acceptable terms. Many of these risks and uncertainties are beyond our control and depend on the actions of third parties. Further information about these and other risks is included in the “Risk Factors” and “Operating and Financial Review and Prospects” sections of 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, 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 obligation to update any forward-looking statements, except as required by law.

Statements that “we believe” and similar expressions reflect our beliefs and opinions, based upon information available as of the date of this presentation. Such statements are inherently uncertain, and investors are cautioned not to place undue reliance on them.

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 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 net 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 include 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.

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.

Amortization of debt discount and issuance costs

We also adjust net income/(loss) for interest expense representing amortization of the debt discount and issuance costs related to our convertible senior notes due 2029 and 2031 issued in Q2 2025. 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

