November 12, 2026

# Nebius Group
# Letter to shareholders Q3 2026

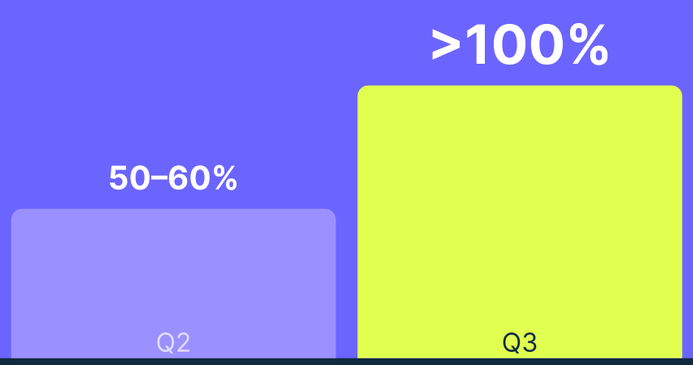
Q3 2026

NEBIUS

## Better payback and prepayments

### Capex coverage by prepayment[1]

Share of capex covered by prepayments excluding deals less than 1 year




### Payback excluding prepayments

Expected payback on a revenue basis for 1–3-year comparable deals




**Pre-tax breakeven by the end of 2027**

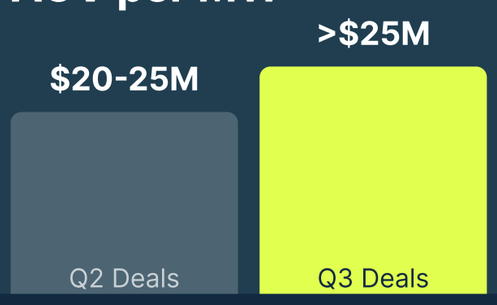
## Customers: a broad and expanding base

### New business generated

**$12B TCV**

Signed since the beginning of Q3

### ACV per MW[2][3]




### A broader base in Q3

**Enterprise AI leaders joined frontier labs and AI-native companies on the platform**

databricks · DATADOG · snowflake

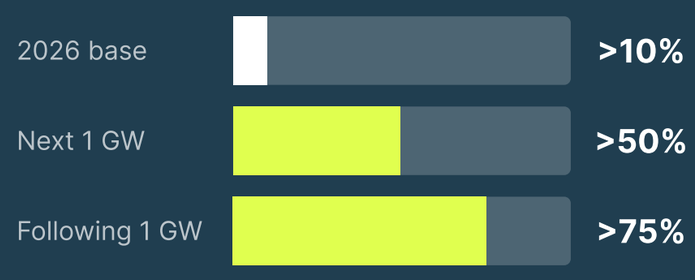
## Our ecosystem

TD SYNNEX · CDW · Reflection · NVIDIA · OpenHands · LangChain · World Wide Technology · Hugging Face · Palantir

## Scaling capacity




### Majority of new capacity to come from owned data centers




## Financial highlights and guidance

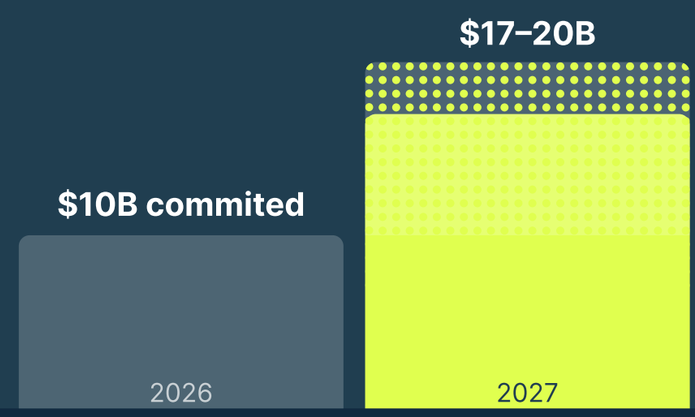
Nebius Group revenue — Q3

**Up 571% YoY to $981M**

Nebius AI Cloud — Q3

**53% adj. EBITDA margin**

Customer prepayment expectations




2027 preliminary revenue guidance




2027 capex



**75%-85%**

covered by prepayments and available investment grade asset backed financing

## Other businesses and investments

Strategic equity stakes — ClickHouse · Toloka

Other businesses — AVRIDE · tripleten

(1) Capex coverage calculated from deals one year or longer, including GPU, data-center construction and leasehold improvement costs
(2) ACV — Annual contract value.
(3) Weighted avg. of deals up to 5 years

# CEO’s letter

**Prepayments on Q3 agreements covered more than 100% of the associated capex. We expect to reach pre-tax breakeven by the end of 2027.**



Arkady Volozh,
Founder and CEO

**Dear shareholders,**

The third quarter demonstrated once again the strength of the business we are building. **Demand remains exceptionally strong, and we have the right mechanisms in place to capture it.**

We do not just build and operate data centers and compute, like neoclouds. Nor do we only provide cloud services on top of rented compute, like inference providers.

Nebius does both: we design, build and run the infrastructure, the cloud software, and the managed AI services that run on it. This is a hyperscaler model, engineered as one system for every compute-intensive AI workload – from training and post-training to inference, agentic search and execution. That lets us focus on how we capture the growth in front of us.

**In the third quarter, this translated into our strongest deal terms to date.**

- Our customers committed to longer deals, at higher prices, with more up front. Prepayments on Q3 agreements covered more than 100% of the associated capex, including data center construction and networking, up from 50-60% in Q2. This has continued rising on agreements signed so far in Q4.
- The average payback term improved to 1 year and 7 months, from 1 year and 10 months in Q2, driven by higher contracted prices. This covers both opex and capex.[(1)]
- Token Factory, our managed inference platform, is approaching a multi-hundred-million-dollar run rate, with contract value growing 17x quarter-on-quarter in Q3. This is only the beginning.
- We are reaching a broader base of customers. AI natives are expanding with us, and we have signed our first agreements with leading enterprise data and AI platforms. We reach them directly, and with partners who sell with us.
- The market is telling us what our capacity is worth, through new pricing mechanisms including reserve auctions and live spot capacity prices. This gives customers more ways to get started on Nebius, and gives us real-time insights into demand.

**Our model is underpinned by capacity, delivered across a flexible and diverse global footprint of sites, partners and power.**

- We prefer to build our own sites, which give us greater control over execution and better economics. These sites are on time or ahead of schedule, and we expect they will make up most of the gigawatts we bring to market in the future.
- Where we rely on third parties, there have been delays. This includes in New Jersey, where we are transitioning the power source to Bloom Energy’s sustainable fuel cells. The initial generation was installed, and we are now scaling the capacity.

(1) On comparable one-to-three-year agreements, and assuming five-year depreciation

# CEO’s letter

In recent weeks our industry has been asked questions about how the AI build-out will be financed.

Our focus remains on building a healthy, profitable business for the long term. We will continue to invest across the business – in our capacity and infrastructure build-out, in our platform and products, and in the team that takes this to market.

What this gives us is the ability to command premium prices and terms, combined with lower costs of our infrastructure as a result of designing and building it ourselves, which in turn results in higher margins.

This approach is paying off, and we are able to finance it in a very prudent way, through customer prepayments and a healthy balance sheet.

- We expect prepayments and financing against existing contracts with investment-grade customers to cover 75%-85% of our 2027 capex, and have a wide range of other funding options available.
- Behind that sits a low cost of debt and a healthy balance sheet, which includes cash and our strategic equity stakes.
- We will remain opportunistic around raising additional external capital when it is cost-effective to do so.

**Today we are updating our guidance for 2026:**

- With a portion of this year’s revenue moving into the first part of next year, we now expect 2026 revenue at the low end of our $3.0-3.4 billion range, with an adjusted EBITDA margin of ca. 40%.
- We now expect the ARR we guided to for December to arrive in February, even before New Jersey ramps up to full capacity.

**And we are also sharing our initial outlook for 2027:**

- We expect to reach pre-tax breakeven by the end of 2027, on the back of the structural benefits discussed above.
- We expect revenue of $12-13 billion, reflecting a conservative approach to the pace of our build-out. We continue to expect to bring online more than 1 GW of new capacity each year.
- We anticipate capex of $40-45 billion, with a disciplined and diversified funding mix.
- As the business matures, we are moving to metrics that better reflect its underlying performance. From 2027, we will provide guidance for revenue and pre-tax earnings.

Our confidence in the fundamentals of what we are building has not changed, and the opportunity ahead of us is as large as it has ever been. We will keep going.

**Arkady Volozh,**
**Founder and CEO**

# Customers: Deeper commitments, broader base

**Demand is not the question**

Since the start of Q3, our customers have committed over $12 billion in contract value to Nebius, on our strongest terms to date. And we still have a lot more to sell.

We choose how we sell our capacity and on what terms. This quarter, we also gave customers more ways to buy.

Even with more choice, customers are meeting our terms: bigger deals, for longer, at higher prices, and with more paid up front.

We saw this clearly in Q3 and thus far in Q4.

**Prepayments: Customer prepayments now cover capex and more.**

- Prepayments on deals signed in Q3 covered over 100% of associated capital expenditures, up from 50%-60% in Q2. This trend is continuing into Q4.
- More than 90% of contract value signed in the third quarter carried a prepayment, and the weighted average prepayment as a share of contract value rose by approximately 20 percentage points quarter-on-quarter.
- Recontracted capacity deals also generated prepayments, driving higher returns on capacity we already own. With that capex already paid, the prepayments help us fund the next build.

**Payback duration: Faster payback and greater visibility**

- Payback period improved to 1 year and 7 months in Q3 from 1 year and 10 months in Q2 for comparable 1-3 year deals, driven primarily by higher contracted pricing on new deals. Payback covers operating costs and capital expenditures and is calculated on a revenue basis, and does not factor in any prepayment benefit.
- We landed five-year customer agreements for Nebius AI Cloud, with ACV/MW higher than in previous quarters. This is where duration, scale, and prepayment meet: the customer reserves capacity, at today's prices, for longer.
- The new five-year AI cloud deals add greater duration diversification to our base and drive visibility into high-margin revenue, as payback is less than two years.

**Demand: Higher prices, from a broader base**

- Annual contract value per MW exceeded $25 million[(1)] in Q3, up from >$20 million on our Q2 deals. Prices in Q3 rose on average more than 20% per GPU-hour versus Q2.
- We introduced new pricing mechanisms to let the market set the benchmark: Reserve auctions cleared more than 15% above our reserve price; Live spot prices saw prices at times exceed our on-demand rate; and Short-term clusters, which are pricing at more than twice the reserved rate.
- We are continuing to expand into new markets where we see demand forming. In the APAC region, for example, since launching in Q1, we have signed multiple customer agreements in Japan, Korea, India, Australia and Singapore. We have also appointed country managers to lead go-to-market in South Korea and Japan.

(1) Weighted avg. of 1 to 5 year deals

# Customers: Deeper commitments, broader base

**Prepayments: Customer prepayments now cover capex and more.**

- **Short-term**, at a premium. New ways to start: reserve auctions, live spot pricing for preemptible capacity, on-demand, and three- to six-month clusters priced at a premium. Both reserve-auction winners were new to Nebius. The first has since expanded its contract value by 10x
- **One to three years**, our core. Seeing expansion from every starting point: customers moving from preemptible to reserved capacity, from one GPU generation to the next, and from Token Factory into dedicated infrastructure.
- **Long-term**, now with AI cloud customers too. The five-year agreements we signed were all expansions by existing customers, on new capacity, adding diversification of duration and supporting our financing. The base is expanding in scale as well: AI cloud agreements reached multi-billion-dollar size in early Q4, just months after our first billion-dollar-plus contracts.

**Adoption: Expanding across segments and routes to market**

Leading enterprise data and AI platforms signed multi-year agreements for the first time:

databricks

**Databricks** selected Nebius AI Cloud following a rigorous evaluation. Their decision demonstrates our ability to support demanding AI R&D workloads.

DATADOG

**Datadog** selected Nebius for AI cloud deployment, citing our purpose-built AI infrastructure, deep technical collaboration and responsive support as the foundation for an infrastructure partnership.

Palantir

**Palantir** named Nebius as a preferred AI cloud partner Palantir customers may leverage with the infrastructure needed to finetune and run inference on their Sovereign AI models.

snowflake®

**Snowflake** signed an agreement to support Snowflake on Nebius AI Cloud. The agreement reflects our continued expansion among leading enterprise technology companies.

# Customers: Deeper commitments, broader base

## Adoption: Expanding across segments and routes to market

AI-natives joined alongside them.

**Cognition,** the company behind Devin, an autonomous AI software engineer, is partnering with Nebius AI Cloud for its first at-scale NVIDIA Vera Rubin clusters to support its compute needs, and accelerating product innovation to meet growing enterprise demand.

**Ortet**, a frontier lab for health, partnered with Nebius for its initial Blackwell cluster, and is set to receive an early Vera Rubin allocation as part of its compute scale-up.

**Cursive**, a London-based AI lab founded by former Google DeepMind researchers, calls Nebius "the ideal compute partner."

**Distil Labs**, a Berlin-based post-training platform, is building specialized models on Nebius AI Cloud for AI agents using synthetic data.

We are building an open ecosystem around Nebius with complementary partners supporting key requirements:

- **Asset-light model partners** enable capital and capacity to build and deliver compute capacity to customers on Nebius architecture and cloud platform, expanding the availability of value-added AI compute globally.
- **Channel and go-to-market partners**, including TD SYNNEX, Palantir, CDW and others, as well as our asset-light model, help us reach and serve more customers.
- **Technology providers** across the AI stack are building with Nebius and contributing tools that customers can use together, from open models and agent frameworks to observability and post-training. Through our builder program, we are working with NVIDIA, Hugging Face, Reflection, OpenHands and LangChain to make that ecosystem useful across more stages of the AI development lifecycle.

# Customers: Deeper commitments, broader base

**Customers are expanding across the platform.**

Most customers come to us for capacity and expand because of the platform.

We run dedicated infrastructure and managed services as one vertically integrated stack. That lets us earn more from each megawatt than a pure-play infrastructure provider, and builds a customer base spread across workload types, GPU generations and contract structures.

**Inference:** Token Factory is approaching a multi-hundred-million run rate less than a year after launch. Many customers are now running multiple workloads on Nebius as a result.

Applied Compute Cursive

**Applied Compute** and **Cursive** chose Token Factory to extend their custom-model work from post-training on Nebius AI Cloud into inference, providing a scalable path to serve enterprise models and deepen customer relationships across the AI lifecycle.

hark

**Hark** selected Token Factory to deliver low-latency inference for Hark Pro's computer-use agent and continue to scale with additional use cases. We are seeing similar adoption from other leading personal AI agents as well.

prosus

**Prosus** uses Token Factory to run open models at high volume across its global lifestyle and e-commerce portfolio, which serves 1 billion users across food delivery, grocery and other services.

**Agentic infrastructure.**
More than 3 million developers now build on Tavily, processing 350 million AI agents requests a month.

monday.com

**monday.com** uses Tavily for real-time web search across its AI products, so agents can research companies, enrich contacts and monitor news inside existing workflows.

OceanAI
Member of the MariApps Group

**Ocean** uses Tavily so its email security agents can verify senders and detect impersonation across one billion investigations a month, cutting investigation time from minutes to seconds.

BELLCYBER

**Bell Cyber** uses Tavily to automate threat research for security operations centers handling up to 100,000 alerts a month, reducing investigations from 30 minutes to under five without adding analysts.

IBM

**IBM Research** selected Tavily, for post-training of the open Granite model family, and expanded the scope for a second time in Q3 2026.

# Capacity: Ramp progressing as we shift to owned sites

## Bringing new capacity online to generate revenue

| | |
|---|---|
| **Capacity ramping** | • In Q3, we delivered our biggest quarter-over-quarter increase in active capacity to date.<br>• Four new sites began generating revenue across the US and Europe, and we delivered additional tranches of capacity in New Jersey in Q3 and through November.<br>• We continue to deliver additional capacity in Q4, with further colocation phases scheduled to be deployed across the US, Europe and Israel. |
| **Next-gen chips** | • NVIDIA Vera Rubin is already shaping our business ahead of deliveries at scale. The first Vera Rubin systems are live and being integrated into Nebius AI Cloud ahead of customer deployments.<br>• We will also be an early deployer of NVIDIA Vera CPUs, complementing our GPU fleet with CPU compute for agentic workloads and data processing. We are seeing high demand for high-performing CPUs, and are well positioned to meet the increasing demand.<br>• Nebius is also the first AI cloud to adopt NVIDIA Groq 3 LPX, boosting inference token generation for NVIDIA Vera Rubin NVL72 in Nebius Token Factory.<br>• We achieved the leading per-GPU MLPerf Inference result on DeepSeek R1 with our Vera Rubin configuration, underpinning customers’ confidence in choosing Nebius. |
| **Future growth** | • We continue to expect to bring online more than 1 GW of new capacity each year. Our outlook takes a prudent approach to the timing of when additional capacity comes online.<br>• Our asset light model is expected to start contributing to our available capacity at a small scale in 2027, after an initial phase closed in Q3. |

| | |
|---|---|
| **Shifts in colocation timelines do not reflect structural change** | • Delivery dates at a number of colocation sites shifted into the first half of 2027, including the remaining New Jersey tranches.<br>• We are confident in being able to bring this capacity online in the near term and are working with our partners to do so.<br>• All affected customers have been kept informed throughout, and their plans are being accommodated. |

# Capacity: Ramp progressing as we shift to owned sites

## Shifts in colocation timelines do not reflect structural change

- In New Jersey, we now expect to have delivered more than half of the capacity contracted with Microsoft by the end of the year, with the remaining tranches now scheduled for 2Q27.

  In late Q3, DataOne started to retire some of the site's temporary generation ahead of the move to permanent power. That decision reset the schedule for the remaining tranches. Microsoft has been part of the planning process.

  We are progressing well with our deployment of Bloom fuel cells as the main power source. Bloom generation is already online and we are now scaling this capacity, expecting to deliver additional capacity over the next few weeks.

**Owned AI factories give us greater control over delivery...**

Our owned sites will account for the majority of the capacity we bring online from 2027 onwards.

- We have secured all the power and components we plan to bring online in 2027, and are aligning construction with each site's scheduled power connection.

  We expect the following new owned sites to deploy capacity in 2027: Missouri, Lappeenranta, our second Mäntsälä site, and Alabama. At scale, these sites represent more than 3 GW of capacity, which we expect to be delivered over the next 3 to 3.5 years.

**... And better margin profile**

- Building our own facilities gives us greater control over our added capacity, from breaking ground to serving customers. Over time, these dynamics show up in our margin profile.
- We have full visibility and management of the entire build. This also gives us control over the permitting and pre-construction phases and pace of our build out.
- We can accelerate near-term development by running site and preconstruction work in parallel, or slow down based on the demand we see.
- Buildings, power infrastructure and fit-out are built out over a number of phases (which can be over a few years).
- Once they are in place, the capital that follows is replacement capex: the compute, depreciated over five years, is the only layer replaced as new GPU generations arrive, and each refresh earns revenue on infrastructure that is already there.

**We build with local communities**

- Our confidence in delivery also reflects how we design our facilities and work with communities and regulators across our U.S. portfolio.
- The US political environment may have shifted, but our responsible by design approach has not. We will continue listening to communities, working transparently, covering power costs, minimizing water use, and generating tangible community benefits.

  As a result of our facility design and proactive engagement efforts, we do not anticipate election impacts on our greenfield or co-location sites, including Pennsylvania.

# Capacity: Ramp progressing as we shift to owned sites

**Laying the foundation for future growth**

- We continue to expect to have 5 GW of contracted power by year-end 2026. We now have more than 4 GW contracted, and have already expanded our contracted footprint with:

  A 300 MW expansion in Missouri, where we now have 1.1 GW site under contract;

  We secured exclusive site rights for a planned 100 MW multi-campus AI data center project in Alberta, Canada, with delivery targeted for 2027 and the potential to scale to several hundred megawatts.

  A new 50 MW colocation in South Carolina, our eighth US site, is expected to come online in 2027;

  Our first sites in the Asia-Pacific region, with two colocations in India expected to come online in 2027.

- We are building across the US and international markets, selecting sites based on customer requirements, power availability, delivery economics and political environment.

- Our asset-light model will also contribute to our ability to expand beyond our own contracted power in a cost-efficient way. We recognized revenue from an initial phase of an asset-light partnership in Q3, and look forward to bringing capacity online from this model in the future.

# AI Cloud: One platform, from compute to production AI

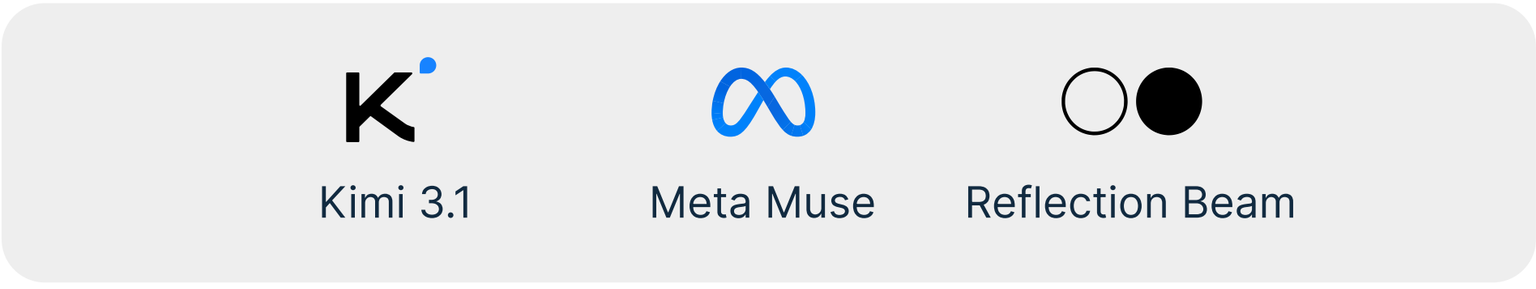
In Q3 we expanded our platform on three fronts: inference, the next workloads (post-training and agents), and enterprise deployment. Each addition gives customers more choice in how they run AI on Nebius, and earns us more of their workload at a lower cost to serve.

Additional proof of the exceptional AI cloud we are building came in September, when we earned Platinum status in SemiAnalysis's ClusterMAX™ 3.0 rating, their highest tier for evaluating AI cloud providers.

## Inference: first to new models, cheaper to serve

Token Factory, our managed inference service, runs leading open-weight and custom models in production and is approaching a multi-hundred-million run rate less than a year after launch. In Q3 we focused on the two things that matter commercially: being first with new models, and lowering what each request costs to serve.

- Kimi K3, GLM 5.3, Reflection Beam and Meta Muse Spark Open were live on Token Factory on their release day, and across the latest DeepSeek, Kimi and GLM models. Token Factory ranks in the top tier of Artificial Analysis benchmarks for performance, latency and quality. Customers can move to a new model the day it lands, on Nebius.




- The Eigen AI and Clarifai integrations were completed in Q3, and in October we added Inferize, which shortens the time required to launch and scale new models.

## Post-training and agents: the next workloads, on the same platform

- Post-training lets customers adapt models to their own data and serve them on Nebius. It sells on infrastructure we already run – the training clusters and Token Factory – which is how Applied Compute and Cursive moved from post-training into inference without leaving the platform.
- Agent execution – sandboxes, agent runtime and model routing – takes us from GPU-intensive training and inference into CPU-intensive workloads. We are bringing CPU capacity online alongside GPUs, and will be an early deployer of NVIDIA Vera CPUs.
- Tavily is the search and data layer for those agents, and brings in customers who do not buy compute from us.

Each layer is sold to the same customer on one platform. That is what lets us generate more revenue from each megawatt than a pure-play infrastructure provider, with a base diversified across workload types.

## Enterprise-ready: connectivity, security and control

Enterprise customers bring requirements that AI natives rarely have: private connectivity to systems that stay where they are, security tooling they already run, and control of cost and capacity across many teams. Those requirements are shaping the core platform, and in Q3 we shipped against all three.

# AI Cloud: One platform, from compute to production AI

## Enterprise-ready: connectivity, security and control

- **Connectivity.** Cloud Interconnect gives a customer a dedicated private connection between Nebius AI Cloud and their own network, so production workloads can run on Nebius alongside systems in on-prem data centers, colocations or other clouds.
- **Security.** Security teams can now assess their Nebius environment with the cloud security posture tools they already use, starting with Wiz — the check that gates enterprise sign-off.
- **Cost and capacity control.** Cost allocation labels, a monthly usage and cost breakdown, and project limits on reserved capacity mean one customer can run many teams on one commitment, with the same visibility over stored data.
- **Speed to production.** Every Nebius service is now available through a REST API with a published OpenAPI specification, and Serverless AI now includes Devlabs persistent workspaces and direct execution of Python jobs — so customers, their CI pipelines and their agents can provision and run on Nebius directly.

A full list of Q3 platform updates can be found here: https://nebius.com/blog/posts/compute-cloud-platform-q3-2026.

## Bringing builders onto the platform

The Nebius AI Builder Program (launched in September), Nebius Academy’s free Agentic AI course co-developed with NVIDIA, and the Nebius x NVIDIA Global AI Hackathon (29,500+ participants) bring the next generation of AI builders onto the platform early, with credits, working examples and access to our engineers.

# Financial update

## Revenue

| In USD $ millions | Three months ended September 30 | | | Nine months ended September 30 | | |
|---|---|---|---|---|---|---|
| | **2025** | **2026** | **Change** | **2025** | **2026** | **Change** |
| Revenues | 146.1 | 980.8 | *571%* | 302.1 | 1,962.1 | *549%* |

**Group revenue in Q3 of $980.8 million was up 571% year-over-year, and up 68% compared to Q2.**

Revenue growth in the quarter was driven by new capacity coming online and revenue from our asset-light model.

- In Q3, we recognized approximately $325 million of non-recurring revenue as part of our asset-light model. This reflects an initial phase of a roll-out, as we engage partners to deploy Nebius architecture and the Nebius AI Cloud platform on their funded capacity.
- We are building a strong pipeline of partnership opportunities. As our asset-light footprint grows, we expect revenue from these activities to become a more regular part of our business.
- Under these partnerships, we earn revenue both when we enable our architecture and platform on a partner’s capacity, and as that capacity comes online.
- Nebius AI cloud revenue was $973.8 million in Q3’26, a 643% increase year-over-year from Q3’25. Our Nebius AI cloud business accounted for approximately 99% of total group revenue during the quarter.
- Annualized run-rate revenue (ARR) of $6.4 billion as of the end of September was up 1,062% year-over-year and up 113% from the $3.0 billion reported at the end of June.

**We no longer see ARR as the right measure of our business. As our business scales, this metric does not reflect the full scope of our base:**

It is based on one month of the quarter, while we do not build our business around any one month;

As we scale new partnerships, this revenue can add choppiness to ARR until such revenue ramps; and

We are also shifting to the next phase, where our business is based on revenue and profitability.

As we are growing and getting mature, we believe revenues and pre-tax earnings are the most appropriate way of evaluating us.

- Remaining performance obligations (RPO) were $38.7 billion as of the end of September, up 3% from $37.5 billion reported at the end of June.

# Financial update

## Operating expenses

| In USD $ millions | Three months ended September 30 | | | Nine months ended September 30 | | |
|---|---|---|---|---|---|---|
| **Expense category** | **2025** | **2026** | **Change** | **2025** | **2026** | **Change** |
| Cost of revenues | 42.9 | 219.7 | *412%* | 97.7 | 457.1 | *368%* |
| as a percentage of revenues | 29% | 22% | | 32% | 23% | |
| Product development | 44.9 | 118.6 | *164%* | 124.2 | 377.0 | *204%* |
| as a percentage of revenues | 31% | 12% | | 41% | 19% | |
| Sales, general and administrative | 89.5 | 237.4 | *165%* | 218.6 | 555.1 | *154%* |
| as a percentage of revenues | 61% | 24% | | 72% | 28% | |
| Depreciation and amortization | 99.0 | 354.7 | *258%* | 223.3 | 826.4 | *270%* |
| as a percentage of revenues | 68% | 36% | | 74% | 42% | |
| **Total operating costs and expenses** | **276.3** | **930.4** | **237%** | **663.8** | **2,215.6** | **234%** |
| as a percentage of revenues | 189% | 95% | | 220% | 113% | |

**Cost of revenue was $219.7 million in Q3'26, representing an increase of 412% 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 operating lease agreements as well as hiring to support our growing operations.
- As a percentage of revenue, cost of revenue was 22% in Q3'26, down from 29% in Q3'25, primarily reflecting operating leverage as we scaled capacity.

**Product development expenses were $118.6 million in Q3'26, representing an increase of 164% 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, including share-based compensation expense.
- As a percentage of revenue, product development expenses decreased to 12% from 31% in the prior year.

**Sales, general and administrative expenses ("SG&A") were $237.4 million in Q3'26, representing an increase of 165% compared to $89.5 million in the same period in 2025.**

- The increase was primarily driven by hiring to support the growth of our business and an increase in outsourced services, as well as consulting, legal and other professional fees primarily incurred in connection with financing transactions.
- SG&A as a percentage of revenues declined to 24% from 61% in Q3'25.

**Depreciation and amortization expenses ("D&A") were $354.7 million in Q3'26, representing an increase of 258% compared to the same period in 2025.**

- The primary driver of the dollar increase in D&A expenses was our continued investment in GPU-related capital expenditures and related data center hardware for the Nebius AI cloud business.
- We use a five-year useful life for our server and network equipment based on usage patterns and current utilization commitments, up from the four-year useful life used prior to 2026.
- D&A as a percentage of revenue decreased to 36% from 68%.

# Financial update

## Adjusted EBITDA

| In USD $ millions | Three months ended September 30 | | | Nine months ended September 30 | | |
|---|---|---|---|---|---|---|
| | **2025** | **2026** | **Change** | **2025** | **2026** | **Change** |
| Adjusted EBITDA / (loss) | (5.2) | 465.8 | *n/m* | (79.9) | 831.5 | *n/m* |
| as a percentage of revenues | -4% | 47% | | -26% | 42% | |

**Group adjusted EBITDA of $465.8 million in Q3 improved by $471.0 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 $518.2 million and adjusted EBITDA margin of 53.2%.
- Compared to Q2’26, group adjusted EBITDA improved by $229.6 million, from $236.2 million to $465.8 million.

## Capital expenditures

- Capital expenditures were approximately $7.3 billion, primarily driven by purchases of GPUs and GPU-related hardware, and data-center expansion activities.

## Capital requirements

We fund our growth from many sources, and we choose among them based on: the cost of capital, dilution, and a healthy balance sheet.

- Prepayments and asset-backed financing of our existing investment-grade contracts alone are expected to fund 75%–85% of our $40–45 billion of capex in 2027. As we choose the mix of contracts we sign — by customer, by duration and by the prepayment attached — we also choose how much external capital we need and when. That is what lets us be opportunistic rather than dependent: we raise when the cost and the timing are right, and are not reliant on any single route to fund 2027.
- Prepayments are now the standard structure of our agreements. We expect $10 billion in committed customer prepayments in 2026, rising to $17–20 billion in 2027. We receive them well in advance, which gives us visibility into our funding. The 2027 range is not a ceiling. It depends on how we allocate capacity across contract types, and we set that mix for the best balance of price, duration and prepayment.
- Our long-term agreements with investment-grade customers can be financed directly. In July we closed the first facility under our asset-backed financing framework: $775 million at SOFR plus 2.50%, secured against long-term agreements with an investment-grade customer. Our agreements with Microsoft and Meta can be financed the same way, and with more than $40 billion of similar committed revenue we see this as a repeatable framework.
- Our balance sheet is healthy: $7.4 billion in cash and cash equivalents at quarter-end, the equity stakes we hold in other businesses, and cash from operations, which is a growing funding source as margins expand.

# Financial update

## Capital requirements

- We have the flexibility to be opportunistic with external capital when it is cost-effective to do so. We will optimize external capital, and any financing we take, against the cost of capital, and balance it across debt and equity. This year that has meant $5.75 billion of convertible notes in August, approximately $1.0 billion of net proceeds under our at-the-market program in the third quarter — a flexible tool, not a commitment — and short-term facilities on attractive terms, such as the $3.5 billion senior secured facility with Citi in October.
- Which sources we draw on, and when, is our decision: further prepayments, additional asset-backed financing, our balance sheet, corporate-level debt or equity. We will continue to evaluate all of them, and pursue the vehicles that offer the most attractive terms and best serve the long-term interests of the business.

## 2026 Guidance

We are updating our 2026 guidance for one change: some third-party capacity we expected to come online this year is now scheduled for the first half of 2027. The contracts behind that capacity remain in place.

- **Revenue:** We now expect revenue at the low end of our $3.0-3.4 billion range.
- **Adjusted EBITDA margin:** Unchanged, at approximately 40%.
- **Capital expenditures:** Approximately $27 billion, above our $20-25 billion range, as we invest in capacity for 2027, in particular at our owned sites.
- **ARR:** We had guided to $7-9 billion by December. We now expect ARR to exceed $7 billion for the month of February, a shift of about two months, even before New Jersey ramps up to full capacity.
- **Contracted power:** We continue to expect 5 GW of contracted power by year-end, from more than 4 GW today.

## 2027: initial outlook

- **Revenue:** $12-13 billion, approximately four times our 2026 revenue.
- **Pre-tax profit:** breakeven in Q4 2027. This is calculated after financing costs, and assumes a five-year depreciation life for compute.
- **Capital expenditures:** $40-$45 billion, sized for more than 1 GW of new capacity. We expect to cover 75-85% of our capex through advanced prepayments and asset-backed financing investment-grade customers along with our balance sheet.
- **Power:** we continue to expect to bring online more than 1 GW of new capacity each year.
- **ARR:** From 2027, we will provide guidance for revenue, pre-tax earnings and capacity. ARR annualizes a single month, and does not reflect the scale of our business.

# Other businesses

## Equity stakes

### AVRIDE

**Autonomous vehicle operations continued to scale**

- The autonomous vehicle fleet grew to 250 vehicles, with a target of 350 by year-end 2026.
- Total autonomous mileage exceeded 2 million miles (3.2 million km).
- Robotaxis have completed more than 100,000 passenger rides. The Dallas operating area has tripled since launch to 31 square miles (80 square km), with further expansion planned.
- Avride plans to begin fully driverless robotaxi operations in the first half of 2027.

**Robodelivery operations at 25 campuses and in more than 30 US cities**

- Avride added 20+ campuses this fall, bringing its total to 25 with more campuses lined up for launch later in 2026 and 2027, including Ohio State, University of Tennessee and University of Kentucky.
- 6,500+ daily deliveries in Q3, up more than 3x year-over-year; cumulative deliveries exceeded 800,000 by the end of Q3.
- Avride expanded its partnership with Wonder on Grubhub, adding 100+ merchants across Jersey City and Hoboken, NJ; Arlington County, VA; and Philadelphia, PA.
- Robodelivery is live in 30+ US markets and in Tokyo, Japan.
- Total robot delivery fleet reached 1,000 units.

### tripleten

**Our edtech business continues to improve efficiency and develop AI offerings**

- Operational efficiencies and tighter cost discipline helped the consumer business cut its cash burn by 44% quarter-over-quarter, and the team remains focused on reaching cash breakeven within B2C.
- Our enterprise-focused B2B offerings have a strong backlog, with courses supporting AI implementation and enablement.
- LatAm achieved positive adjusted EBITDA and adjusted cash flow in every month of Q3.
- In the US, TripleTen is focused on its profitable core programs and repositioned its offerings toward experienced professionals moving into AI roles.

# Other businesses

# Equity stakes

## Toloka

**Toloka has grown from a data provider to a full-cycle partner for AI labs and enterprises**

- In Q3, Toloka continued to grow revenue, doubling quarter-over-quarter, with a healthy and expanding contribution margin as well as continued leverage within the core business.
- Toloka expanded its Arena agentic capabilities evaluation to 35+ models tested on private, non-contaminated agentic tasks. The failure modes Arena surfaces feed directly into Toloka's off-the-shelf data catalog, which continues to expand to meet strong demand from frontier labs. Built once and licensed many times, these ready-to-buy datasets create a highly scalable revenue stream.
- Self-service is now enabled across human and synthetic data collection, labelling, and evaluation, and is already integrated with the Nebius AI Builders Program.
- The post-training offering launched in Q3 and brings the model improvement loop proven with frontier labs to companies like Shopify. For labs, Toloka's ablation studies show that a single epoch of post-training on its OTS dataset improves performance on unseen external benchmarks.

## 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 $400 million in a Series D financing at a valuation of approximately $15 billion. Following this financing, Nebius Group continues to own a significant minority equity stake in ClickHouse.

# Forward-looking statements

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

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

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

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

**Disclaimer**

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

# Use of Non-GAAP financial measures

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

- Adjusted EBITDA / (loss) means U.S. GAAP net income / (loss) from continuing operations before (1) depreciation and amortization, (2) share-based compensation expense, (3) acquisition and other corporate transaction-related costs, (4) interest income, (5) interest expense, (6) income / (loss) from equity method investments, (7) gain from revaluation of investments in equity securities, (8) other income / (loss), net, (9) induced conversion expense, (10) income tax expense / (benefit).
- Adjusted net income / (loss) means U.S. GAAP net income / (loss) from continuing operations before (1) share-based compensation expense, (2) foreign currency exchange gains / (losses), (3) acquisition and other corporate transaction-related costs, (4) amortization of debt discount (accretion) and issuance costs, net of interest expense capitalized, (5) gain from revaluation of investments in equity securities, and (6) induced conversion expense. 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.

# Use of Non-GAAP financial measures

## Share-based compensation expense

Share-based compensation expense 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 currency exchange gains / (losses)

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

## Acquisition and other corporate transaction-related costs

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

## Amortization of debt discount (accretion) 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 difference between the amount payable at maturity and the initial carrying amount of the notes, and is amortized to interest expense over the term of the notes, 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.

## Induced conversion expense

Induced conversion expense reflects the loss recognized on conversion of our convertible notes, including induced conversions in which the fair value of the consideration issued in excess of the fair value of the securities issuable under the original conversion terms.

NEBIUS