

METIS RESEARCH

AI Infrastructure / Neocloud

Nebius Group N.V.

NASDAQ: NBIS

Exceptional Growth, Fully Priced: Initiating at Neutral


Prepared: July 27, 2026

Investment Summary

We initiate coverage of Nebius Group (NASDAQ: NBIS) with a NEUTRAL (Hold) rating and a 12-month price target of \$198, essentially level with the July 27, 2026 close of \$198.25. Nebius is one of the highest-quality, best-financed independent AI-infrastructure platforms in the market FY2025 revenue rose 479% (on the restated, Toloka-excluded base) and Q1 2026 revenue grew 684% YoY, the core AI-cloud business targets \$7–9B of exit-2026 ARR, and NVIDIA has taken a ~9.3% equity stake. After a 262% advance over the past twelve months, however, the shares already discount a great deal of that success, and our base-case fair value sits almost exactly at the current price.

Metric	Value	Metric	Value
Rating	NEUTRAL (Hold)	52-week range	\$50.00 – \$299.86
Current price (Jul 27, 2026)	\$181	Market cap	~\$58.7B
12-month price target	\$198	Enterprise value	~\$57.3B
Implied upside	~8.5%	Diluted shares	~296M
FY2026E revenue	\$3.0–3.4B (guidance)	Net cash	~\$1.5B
FY2026E exit ARR	\$7–9B	Beta	~1.4

Figure 1 — NBIS Share Price Performance (Jul 2025 - Jul 2026)



Source: Public market data (approximate monthly closes); NBIS relisted Oct 2024. Prepared Jul 27, 2026.

Figure 1 NBIS share price, Jul 2025–Jul 2026 (approximate monthly closes). Source: public market data.

Key investment points

- **Hyper-growth backed by \$40B+ of contracted revenue, but supply-gated.** Revenue grew 684% YoY in Q1 2026 to \$399M and management guides FY2026 revenue of \$3.0–3.4B with \$7–9B of exit ARR. Microsoft (up to \$19.4B) and Meta (up to ~\$27B) anchor a backlog exceeding \$40B, so growth is limited by how fast Nebius can build capacity not by demand. The company sells essentially all capacity it brings online.
- **Full-stack differentiation and a preferential NVIDIA relationship.** Unlike scale-focused peers, Nebius layers Token Factory (managed inference), agentic search (Tavily) and inference optimization (Eigen AI) on top of raw GPU compute. NVIDIA's \$2B investment (~9.3% stake) and a joint plan toward ~5 GW by 2030 signal preferential access to the scarcest silicon (including the forthcoming Vera Rubin platform).

- **Strong unit economics but GAAP losses through the build-out.** Core AI-cloud adjusted EBITDA margin was already 45% in Q1 2026, and we model group adjusted EBITDA margin rising from ~40% (FY2026E) to ~50% (FY2030E). Yet D&A consuming roughly half of revenue and rising interest keep GAAP earnings negative through 2030 on our conservative base case (which excludes lumpy equity-stake gains).
- **Valuation is full and almost entirely terminal-dependent hence Neutral.** Our DCF base case is \$205/share and trading comps \$189; the blended \$198 target sits right at today's price. Because explicit-period free cash flow is deeply negative by design, more than 100% of DCF value is terminal a fragile setup. We would turn constructive on a pullback toward the \$150s or on evidence the asset-light pivot structurally lifts free cash flow.

Summary financials

\$M unless noted	FY2024A	FY2025A	FY2026E	FY2027E	FY2028E	FY2029E	FY2030E
Revenue	92	530	3,200	7,500	12,000	16,500	21,000
YoY growth	—	479%	504%	134%	60%	38%	27%
Gross profit	48	363	2,176	5,175	8,400	11,550	14,700
Adjusted EBITDA	(226)	(65)	1,280	3,300	5,640	8,085	10,500
Adj. EBITDA margin	(247%)	(12%)	40%	44%	47%	49%	50%
EBIT	(400)	(596)	(576)	(660)	(192)	578	1,596
Net income	(352)	29	(612)	(1,397)	(1,506)	(1,041)	(132)
Diluted EPS (\$)	(1.49)	0.11	(2.23)	(4.90)	(5.19)	(3.54)	(0.44)
Free cash flow	(1,077)	(3,664)	(18,250)	(14,253)	(7,519)	(2,867)	(61)
Exit ARR	—	1,250	8,000	13,500	18,000	22,000	26,000

Source: rebuilt model on reported filings. A = actual (FY2024 restated ex-Toloka), E = estimate. GAAP net income excludes lumpy equity-stake revaluation gains in projections (conservative).


Investment Thesis

Our Neutral rating balances an exceptional business against a full valuation. Below we lay out the four pillars that frame the investment case, then explain why despite genuine admiration for the asset we do not see enough margin of safety at current levels to rate the shares a Buy.

Pillar 1 A supply-gated, \$40B+ contracted growth engine

Nebius has moved from a single Finnish data center to a multi-gigawatt global build-out in roughly six quarters. Revenue compounded from \$91.5M (FY2024A, restated to exclude Toloka) to \$529.8M (FY2025A, +479%) and reached \$399M in Q1 2026 alone (+684% YoY, +75% QoQ). The forward book is even more striking: the Microsoft agreement (up to \$19.4B through 2031) and the expanded Meta agreement (up to ~\$27B from 2027) give Nebius a contracted backlog exceeding \$40B extraordinary for a company with under \$1B of trailing revenue. Management repeatedly emphasizes that demand exceeds supply: pricing is firm, contract durations are lengthening, and capacity sells out. The binding constraint is execution data-center construction, power availability and GPU delivery not sales. This reframes the key question from “can they sell it?” to “can they build and deliver on time?”

Figure 2 — Revenue Growth Trajectory (2024A-2030E)




Source: Company filings; Task 2/3 model & valuation. Prepared Jul 27, 2026.

Figure 2 Revenue growth trajectory, 2024A–2030E.


Pillar 2 Full-stack differentiation and NVIDIA validation

Whereas the largest competitor, CoreWeave, has leaned into raw scale, Nebius is building a fuller software stack on top of its compute: Nebius Token Factory (managed inference across major open-source models), agentic search (Tavily, acquired ~\$275M), and frontier inference optimization (Eigen AI, ~\$643M). The bet is that enterprises want an integrated platform that shortens the distance from model to production, not commodity GPU rental. NVIDIA’s \$2B equity investment (~9.3% stake) and a joint build-out plan toward as much as 5 GW by 2030 are powerful signals of preferential access to scarce GPUs and next-generation platforms (Vera Rubin). The trade-off is deepening dependence on a single silicon supplier that is also a strategic investor across the sector.

Figure 19 — ARR & Contracted Power Build-out*Figure 19 Core AI-cloud exit ARR and contracted power build-out.*

Pillar 3 Powerful operating leverage as capacity monetizes

The core AI-cloud business already runs at a 45% adjusted EBITDA margin (Q1 2026), and we model group adjusted EBITDA margin rising from ~40% in FY2026E (management guidance) to ~50% by FY2030E as fixed costs are spread across a much larger revenue base. Gross margin (excluding D&A) holds around 68–70%. The catch and a central reason for our Neutral is that depreciation on the rapidly deployed GPU fleet consumes roughly half of revenue, so GAAP operating income stays negative until 2029E and net income remains negative through 2030E on our conservative base case (which assigns zero value to the lumpy equity-stake gains that turned FY2025 GAAP net income positive). Adjusted EBITDA and ARR are the cleaner read-throughs on the underlying economics.


Figure 11 — Adjusted EBITDA Margin Progression

Source: Company filings; Task 2/3 model & valuation. Prepared Jul 27, 2026.

Figure 11 Adjusted EBITDA margin progression, 2024A–2030E.

Pillar 4 A quality asset, but the price already reflects it

Our blended 12-month target of \$198 (DCF 60% at \$205, trading comps 40% at \$189) sits essentially on top of the current \$198.25. On EV/2027E revenue, Nebius trades at 7.6x versus a neocloud peer median near 4.2x a ~2x premium that we believe is partly, but not fully, justified by ~5x faster growth, a cleaner balance sheet, and full-stack differentiation. Critically, because explicit-period free cash flow is deeply negative during the build-out, more than 100% of our DCF value is terminal value; the fair value is therefore highly sensitive to the exit multiple and WACC. Combined with live customer-concentration risk (Microsoft and Meta) and a 2026–2028 sector refinancing wall, this leaves a balanced risk/reward at today's price a Hold, not a Buy.

Figure 32 — Valuation Football Field (KEY CHART)

Source: Company filings; Task 2/3 model & valuation. Prepared Jul 27, 2026.

Figure 32 Valuation football field.

Risk Factors

We view the following as the most material risks to our thesis and price target, organized by category. The overarching risk is that Nebius is valued almost entirely on cash flows that arrive beyond our explicit forecast horizon, making the equity unusually sensitive to execution, financing and terminal assumptions.

Company-Specific Risks

- **Customer concentration.** A large share of Nebius's contracted backlog rests on two counterparties Microsoft (up to \$19.4B) and Meta (up to ~\$27B). While both are investment-grade, the loss, delay, or renegotiation of either contract would materially affect revenue visibility and the financing that depends on those prepayments. Diversifying the customer base and successfully re-selling discretionary Meta tranches to third parties are critical mitigants.
- **Execution and delivery risk.** Growth is gated by the ability to build data centers, secure power, and take GPU delivery on schedule. Contracts (e.g., Microsoft's) contain delivery-date provisions; failure to meet agreed dates, absent alternative capacity, can give customers termination rights. Constructing multiple gigawatt-scale facilities on time and on budget is operationally demanding and largely unproven at the targeted 2030 scale.
- **Key-person dependency.** The company is closely identified with founder-CEO Arkady Volozh, who also controls a large economic stake. His vision, relationships (including with NVIDIA), and capital-allocation judgment are central to the thesis; an unexpected departure or incapacity would be difficult to replace and could weigh on strategy and sentiment.
- **Technology obsolescence.** GPU fleets depreciate rapidly against each new NVIDIA generation (Hopper → Blackwell → Vera Rubin). Depreciation already consumes roughly half of revenue, and mistimed purchasing or slower-than-expected utilization of older hardware could erode returns on very large capital outlays.
- **Integration and M&A risk.** Rapid inorganic expansion (Tavily, Eigen AI, Clarifai, and reported interest in others) carries integration, retention, and valuation risk e.g., \$643M for Eigen AI's ~20-person team implies a high per-employee price that depends on retaining and scaling that talent.

Industry / Market Risks

- **Vertical integration by customers.** The neoclouds' largest customers are also their most capable potential competitors. Reports in mid-2026 that Meta was exploring a "Meta Compute" offering briefly cut ~15% from Nebius and CoreWeave shares, illustrating how hyperscaler self-build or excess-capacity resale could compress third-party demand and pricing.
- **Competitive intensity.** CoreWeave is roughly ten times larger by GPU fleet and better-capitalized in backlog, while numerous well-funded neoclouds (Crusoe, Lambda, Nscale, Applied Digital) and the hyperscalers themselves compete for the same GPUs, power, and customers, pressuring pricing and share over time.
- **Commoditization of GPU rental.** Raw GPU-as-a-service risks becoming a commodity; Nebius's software-stack strategy (Token Factory, agentic search) is a deliberate hedge, but if differentiation fails to stick, margins on core compute could compress.
- **Regulatory, power, and geopolitical factors.** Advanced-chip export controls, data-center permitting and energy constraints, sustainability scrutiny, and the company's Yandex heritage (now fully separated from Russia) all present regulatory and reputational considerations that could affect where and how quickly it can build.

Financial Risks

- **Capital intensity and GAAP unprofitability.** Guided 2026 capex of \$20–25B dwarfs current revenue, and no major neocloud is yet GAAP-profitable given heavy depreciation. Reported net income is further distorted by


large non-cash valuation swings on non-core stakes (a \$621M Q1 2026 gain versus a \$250M Q4 2025 loss), complicating the earnings picture.

- **Financing and refinancing risk.** The build-out relies on prepayments, \$4.3B of convertibles, NVIDIA equity, and asset-backed debt (total debt approaching ~\$8B). Analysts flag a 2026–2028 “refinancing wall” of GPU-collateralized loans across the cohort, concentrated among a small circle of lenders a systemic risk if credit conditions tighten.
- **Shareholder dilution.** Share count rose roughly 20% over the past year, and continued equity issuance (including shares used in acquisitions and potential future raises) could dilute existing holders, even as management prioritizes non-dilutive funding where possible.

Macroeconomic Risks

- **AI-spending cyclicality.** The thesis assumes sustained, multi-year hyperscaler and enterprise AI investment. A slowdown in AI capex whether from disappointing model economics, macro weakness, or an “AI winter” would hit demand, pricing, and financing simultaneously, given how correlated the sector is.
- **Interest-rate and financing-cost sensitivity.** As a capital-intensive, debt-financed grower, Nebius is sensitive to interest rates and credit spreads; higher rates raise the cost of the debt underpinning the build-out and pressure the valuations of long-duration growth equities such as NBIS (beta above 3).
- **Currency and geographic exposure.** Reporting in U.S. dollars while operating across Europe, the Middle East, and the U.S. creates FX translation exposure, and geographic concentration of power and construction subjects the company to regional energy, regulatory, and geopolitical shifts.

Figure 18 — Competitive Benchmarking: Nebius vs CoreWeave



Source: Company filings; Task 2/3 model & valuation. Prepared Jul 27, 2026.

Figure 18 Competitive benchmarking: Nebius vs CoreWeave.

Company Description

Nebius Group N.V. is an Amsterdam-headquartered technology company that designs, builds, and operates full-stack artificial-intelligence (AI) cloud infrastructure. It sits within the emerging “neocloud” category specialist providers that rent high-performance GPU compute to AI developers, enterprises, and hyperscalers, positioning themselves as focused alternatives to the general-purpose hyperscale clouds (Amazon Web Services, Microsoft Azure, and Google Cloud). Nebius describes itself as building “the full-stack platform for developers and companies to take charge of their AI future from data and model training to production deployment.” In practice this means it owns and operates large-scale GPU clusters, layers a purpose-built cloud software stack on top, and increasingly offers managed inference, agentic, and developer tooling above the raw compute.

The company makes money principally by selling access to GPU-based compute capacity, generally under multi-year contracts with large, creditworthy customers, supplemented by shorter-term, on-demand and reserved cloud consumption from AI-native startups and digital enterprises. Revenue is recognized as capacity is delivered and consumed; management’s headline forward metric is annualized run-rate revenue (ARR), calculated as the final month’s AI-cloud revenue in a quarter multiplied by twelve. Deal economics are unusually large by cloud-software standards: Nebius has signed a contract with Microsoft valued at up to \$19.4 billion and an infrastructure agreement with Meta worth up to approximately \$27 billion, giving it a contracted revenue backlog exceeding \$40 billion before counting on-demand demand.

Nebius operates a rapidly expanding global data-center footprint. In roughly three quarters it went from a single site in Finland to multiple locations across Europe, the United States, and the Middle East, and it is now developing a series of large “AI factories,” including a 1.2 GW proprietary facility in Pennsylvania, a large campus in the U.S. Midwest, and expansions in Finland and France. The economics of the business are therefore driven less by traditional software metrics and more by physical scale: contracted and connected electrical power (measured in megawatts and gigawatts), GPU fleet size, and the pace at which contracted capacity can be converted into revenue-generating active capacity.

Financially, Nebius is in hyper-growth. Full-year 2025 group revenue was \$529.8 million, up roughly 479% (on the restated, Toloka-excluded basis) year over year, and the core AI-cloud business exited 2025 with \$1.25 billion of ARR ahead of the \$900 million to \$1.1 billion range management had guided. In the first quarter of 2026 the company reported revenue of \$399.0 million, up 684% year over year and 75% sequentially, with the core Nebius AI cloud generating an adjusted EBITDA margin of about 45%. Management guides to full-year 2026 revenue of \$3.0–\$3.4 billion, ARR of \$7–\$9 billion, and a roughly 40% group adjusted EBITDA margin.

It is worth dwelling on the gap between reported revenue and ARR, because it is central to reading the business. Revenue is recognized as capacity is delivered and consumed over a quarter, whereas ARR annualizes only the final month’s AI-cloud revenue so in a company adding capacity every month, ARR runs well ahead of trailing revenue and functions as a forward indicator of the run-rate the business is exiting each quarter at. This is why exit-2025 ARR of \$1.25 billion coexists with full-year 2025 revenue of \$529.8 million, and why management guides 2026 to \$3.0–\$3.4 billion of revenue but \$7–\$9 billion of exit ARR. On unit economics, recent disclosures imply roughly \$0.9 million of ARR per active megawatt of capacity, giving investors a rough yardstick for translating the company’s multi-gigawatt power ambitions into a future revenue base though realized economics will depend on pricing, utilization, and the mix between dedicated hyperscaler contracts and higher-margin on-demand consumption.

A distinctive feature of the company is its portfolio of non-core assets. Beyond the core cloud, Nebius owns the autonomous-mobility company Avride and the ed-tech reskilling platform TripleTen, and it holds significant equity stakes in the open-source analytics database company ClickHouse and in the AI-data company Toloka (which was deconsolidated in 2025 after an outside investment led by Bezos Expeditions). Management has been explicit that these stakes are a strategic funding engine a source of billions of dollars of potential capital to reinvest in the capital-

hungry core cloud without diluting shareholders. This dual structure an aggressively scaling infrastructure business supported by a set of monetizable technology holdings is central to the investment narrative and to how the company finances its build-out.

As of late July 2026, Nebius carried a market capitalization of roughly \$58.7 billion on a share price near \$190, having traded as high as \$299.86 in June 2026 and as low as \$50 over the prior twelve months. The stock is highly volatile (beta above 3), reflecting both the extraordinary growth trajectory and the significant execution, financing, and customer-concentration risks discussed later in this report. In summary, Nebius is best understood as a well-capitalized, founder-led, European-domiciled pure-play on the build-out of AI compute infrastructure, differentiated by an in-house full-stack engineering culture, a close relationship with NVIDIA, and a balance sheet supported by both hyperscaler prepayments and valuable non-core holdings.

Figure 5 — Nebius Group at a Glance

Ticker / Exchange NBIS · NASDAQ	HQ Amsterdam, Netherlands
FY2025 revenue \$529.8m (+351%)	FY2026E revenue \$3.0–3.4bn (guidance)
Exit-2026 ARR target \$7–9bn	Market cap ~\$58.7bn
Employees ~1,371 (>1,000 engineers)	Contracted power >4 GW (~5 GW by 2030)
Anchor customers Microsoft, Meta	Strategic investor NVIDIA (~9.3%)

Source: Company filings; Task 2/3 model & valuation. Prepared Jul 27, 2026.

Figure 5 Nebius at a glance. Source: company filings

Company History

Nebius's corporate lineage runs directly through Yandex, one of the defining technology companies of the Russian-speaking internet. The search engine that became Yandex launched in 1997, and in 2007 a Dutch holding company, Yandex N.V., was established as the group's parent. Yandex N.V. raised approximately \$1.3 billion in a May 2011 initial public offering on NASDAQ, becoming one of the highest-profile technology listings to emerge from the region. Under co-founder and long-time chief executive Arkady Volozh, Yandex grew into a diversified internet conglomerate spanning search, advertising, ride-hailing, e-commerce, cloud, autonomous vehicles, and more.

The company's trajectory was upended by Russia's 2022 invasion of Ukraine. Volozh stepped down as chief executive in 2022, and in 2023 he was placed on the European Union sanctions list a designation later lifted in early 2024. Faced with an untenable position straddling Russian operations and a Western listing, the group pursued a landmark separation: in 2024 Yandex N.V. sold its Russian businesses to a consortium of Russian investors in a transaction valued at roughly \$5.4 billion, one of the largest corporate exits from Russia since the war began. The Russian assets retained the Yandex name and brands; the Dutch-listed parent kept a set of international, non-Russian technology assets and cash proceeds.

That remaining international entity was renamed Nebius Group in 2024, taking its name from an internal AI-cloud initiative that had been incubated within Yandex. The name “Nebius” and the core cloud business trace back to Yandex’s in-house infrastructure and machine-learning teams, many of whom relocated to hubs in Amsterdam, elsewhere in Europe, and Israel. Volozh resumed the chief-executive role in August 2024, and Nebius’s shares which had been suspended from trading on NASDAQ during the sanctions and restructuring period resumed trading under the ticker NBIS in October 2024. The re-listing effectively functioned as a fresh start for a company that, while young as “Nebius,” carried more than two decades of accumulated engineering talent and infrastructure expertise.

From that point the company moved with unusual speed. Through late 2024 and 2025 it rapidly expanded its GPU cloud footprint, opening or announcing data centers across Finland, elsewhere in Europe, the United States, and the Middle East. It cultivated an early customer base of AI-native companies including France’s Mistral, coding platform Cursor, and image-AI company Black Forest Labs before broadening toward larger digital enterprises. The pivotal inflection came in 2025 with a wave of hyperscaler contracts: a five-year agreement with Microsoft announced in September 2025, valued at approximately \$17.4 billion (up to \$19.4 billion with additional services), anchored on a data center in Vineland, New Jersey; followed by an initial Meta agreement in late 2025 and a substantially expanded Meta deal in March 2026 worth up to roughly \$27 billion.

Two other 2025–2026 developments reshaped the company. First, NVIDIA made a \$2 billion equity investment, taking a stake of roughly 9.3% and entering a partnership to jointly develop next-generation infrastructure and to support a build-out toward as much as 5 GW of capacity by 2030 a powerful signal of preferential access to scarce GPUs. Second, the company began actively monetizing and reshaping its non-core portfolio: it deconsolidated Toloka after an investment led by Jeff Bezos’s Bezos Expeditions, and it benefited from sharply higher valuations of its ClickHouse stake, which contributed large non-cash gains to reported earnings.

The most recent chapter, spanning late 2025 into 2026, has been defined by capital formation and inorganic expansion. Nebius raised \$4.3 billion in convertible senior notes and secured its first asset-backed debt facilities (including a \$775 million senior secured facility in July 2026), while pursuing a string of acquisitions to build out the software layer above its compute: agentic-search company Tavily (roughly \$275 million, February 2026), inference-optimization specialist Eigen AI (roughly \$643 million, May 2026), and inference platform Clarifai, among others. Reports have also linked Nebius to interest in Israeli model developer AI21 Labs. In under two years, the company has transformed from a sanctions-era restructuring into one of the most closely watched AI-infrastructure names on NASDAQ.

Figure 6 — Key Milestones



Source: Company filings; Task 2/3 model & valuation. Prepared Jul 27, 2026.

Figure 6 Key milestones. Source: company disclosures.

Management Team

Arkady Volozh Founder, Chief Executive Officer & Executive Director

Arkady Volozh is the defining figure at Nebius: founder, chief executive, executive director, and its largest individual shareholder with an economic interest of roughly 11.5% (worth several billion dollars at recent prices). A serial entrepreneur trained in applied mathematics at the Gubkin Institute of Oil and Gas, Volozh built his career at the intersection of computer science and company-building. Before Yandex he founded and led IT enterprises including InfiNet Wireless and CompTek International, serving as CompTek's chief executive before moving to run Yandex. He co-founded Yandex and served as its chief executive for roughly 25 years, from 2000 to 2022, transforming a search engine into a diversified internet conglomerate and one of the few genuinely global technology companies to emerge from the region.

Volozh stepped down from the Yandex chief-executive role in 2022 amid the geopolitical fallout of the war in Ukraine and was briefly subject to EU sanctions in 2023 before that designation was lifted in early 2024. He resumed the chief-executive role at the renamed Nebius Group in August 2024, personally driving the company's pivot into AI infrastructure. Beyond his operating record, Volozh has an established track record as a technology investor: he backed and sat on the board of Face.com, an Israeli facial-recognition company acquired by Facebook in 2012; was an early investor in the quick-commerce pioneer Getir; and has served on the board of Neurosteer, a medical-grade brain-activity interpretation company. His combination of founder ownership, long operating tenure, and investor sensibility underpins Nebius's owner-operator culture and its explicit strategy of treating non-core stakes as a capital engine for the core business.

Dado Alonso Chief Financial Officer

Dado Alonso serves as chief financial officer, presenting the company's quarterly results and leading its increasingly sophisticated financing strategy. That strategy has become central to the Nebius equity story: rather than fund its multi-billion-dollar data-center build-out primarily through dilutive equity, the company has combined customer prepayments (notably from Microsoft and Meta), \$4.3 billion of low-coupon convertible senior notes (with 1.25% and 2.60% coupons), a \$2 billion strategic equity investment from NVIDIA, and, from mid-2026, asset-backed debt facilities collateralized by GPU infrastructure and long-term contracts. Alonso has framed the goal as financing the bulk of 2026 capital expenditure from internal cash flows and prepayments, converting operational infrastructure assets into growth capital while preserving the balance sheet. His communication emphasizes ARR and adjusted EBITDA as the cleaner read-throughs on performance, given the noise that non-cash valuation adjustments introduce into GAAP net income.

Roman Chernin Co-Founder & Chief Business Officer

Roman Chernin is a co-founder of Nebius and its chief business officer, responsible for commercial strategy, customer relationships, and the structuring of the large hyperscaler contracts that anchor the backlog. Chernin brings roughly a dozen years of experience leading the development of major digital services at Yandex, including Search and Maps, giving him deep product-and-platform credibility alongside commercial responsibility. He has articulated the company's disciplined approach to its landmark deals prioritizing margin and system-level economics over headline revenue, and structuring large contracts so that they help finance the rest of the build-out. He has also been the public voice on the acquisition strategy, describing a "capacity-scarcity world" in which optimized inference and infrastructure scale are the scarce, valuable capabilities the company is assembling through both organic engineering and targeted M&A such as Eigen AI and Tavily.

Andrey Korolenko Chief Product & Infrastructure Officer; Ophir Nave Chief Operating Officer


Andrey Korolenko, chief product and infrastructure officer, leads the technology and infrastructure organization the GPU clusters, data-center engineering, and cloud-software stack that constitute the company's core differentiation.

Ophir Nave, chief operating officer, oversees operations across the group's expanding global footprint. Both draw on the pool of more than one thousand engineers many with backgrounds at Microsoft, Amazon, Yandex, and other leading technology companies that management repeatedly cites as its single greatest competitive asset. The R&D core is anchored in Amsterdam with additional hubs across Europe and Israel, and, following the Eigen AI acquisition, a growing engineering and research presence in the San Francisco Bay Area.

Governance and ownership

Nebius is chaired by John Boynton, and the board also includes Johanna Boynton among its directors. Governance reflects a founder-controlled company: Volozh's roughly 11.5% economic stake, combined with the group's dual-class share structure (the listed securities are Class A shares), concentrates influence with the founder and long-tenured insiders, whose average board and management tenure is around two years given the company's recent reconstitution. Insider activity in 2026 has included routine on-market share sales by Volozh and other insiders (collectively tens of millions of dollars), which some observers monitor as a sentiment signal, though such sales are common for founders diversifying concentrated holdings. On balance, the management team pairs rare operating pedigree with heavy inside ownership a governance profile that aligns incentives with long-term value creation but also concentrates key-person risk in the founder.

Figure 7 — Organizational Structure



Source: Company filings; Task 2/3 model & valuation. Prepared Jul 27, 2026.

Figure 7 Organizational structure. Source: company filings.

Products & Services

Nebius's offering is best understood as a vertically integrated, AI-native cloud stack that spans the full lifecycle of AI development from raw infrastructure through training, inference, and increasingly agentic applications. At the base is the core Nebius AI cloud: large-scale clusters of the latest NVIDIA GPUs (from H100/H200 and Blackwell-class systems toward next-generation Vera Rubin deployments), delivered as bare-metal and multi-tenant cloud capacity with high-performance networking, storage, and orchestration purpose-built for intensive AI training and inference workloads. Because Nebius engineers much of its own hardware integration, data-center design, and cloud software in-house, it markets superior price-performance and GPU utilization versus general-purpose clouds, and

has earned a “Gold” tier placement (alongside Oracle and Azure, below CoreWeave’s “Platinum”) in independent neocloud rankings.


Capacity is sold in two broad modes. The first is large, dedicated, multi-year capacity contracts with hyperscalers and frontier labs the Microsoft (up to \$19.4 billion) and Meta (up to ~\$27 billion) agreements are the flagship examples, delivered in tranches across specific data centers and, in Meta’s case, anchored on one of the first large-scale deployments of NVIDIA’s Vera Rubin platform from early 2027. The second is on-demand and reserved cloud consumption by AI-native companies and digital enterprises, priced by GPU-hour and increasingly by tokens processed. Management has emphasized that demand persistently exceeds supply “everything we build, we sell” with strong and rising pricing, longer contract durations, and full capacity sell-outs across recent quarters.

Above the raw compute, Nebius has been rapidly assembling a software and services layer. Nebius Token Factory is its managed inference platform, providing enterprise-grade autoscaling endpoints and fine-tuning pipelines across major open-source models (DeepSeek, GLM, GPT-OSS, Kimi, Llama, MiniMax, Qwen, and others). The May 2026 acquisition of Eigen AI a Bay Area startup founded by MIT HAN Lab researchers and recognized by NVIDIA as a leading speed-inference provider was explicitly aimed at making Token Factory a frontier managed-inference platform, maximizing tokens per GPU and improving unit economics. The February 2026 acquisition of Tavily added agentic search, giving AI agents real-time web access for factual grounding, while the Clarifai acquisition further strengthened inference optimization. Together these position Nebius to serve not just model training but production AI systems and autonomous agents.

A further pillar of the offering is Nebius’s tight alignment with NVIDIA’s reference architectures. The company was among the first cloud providers globally to achieve NVIDIA’s designation for large-scale training clusters (including GB200 NVL72-class systems), and the expanded Meta agreement is structured around one of the first large-scale deployments of the forthcoming Vera Rubin platform. This preferential positioning matters for two reasons: it helps secure allocation of the scarcest GPUs during a period of chronic shortage, and it lets Nebius offer customers early access to each new performance tier a meaningful differentiator when model developers compete on training throughput and inference cost. The trade-off is deepening dependence on a single silicon supplier that is simultaneously a strategic investor and, indirectly through its stakes across the sector, a shaper of the competitive field.

The company also offers AI Studio and developer-facing tooling, MLOps capabilities (including collaborations such as one with Saturn Cloud supporting NVIDIA AI Enterprise), and vertical initiatives such as an AI-for-healthcare and life-sciences program. Distribution increasingly runs through partners: Nebius has integrations with LangChain’s agent frameworks and a go-to-market relationship with TD SYNEX that surfaces dedicated GPU capacity to channel partners and enterprise customers. This full-stack, software-plus-silicon positioning is the company’s central strategic bet: whereas the largest competitor, CoreWeave, has focused primarily on scale, Nebius is wagering that enterprise customers want both scale and a richer software stack that reduces the gap between building a model and running a real product.

Rounding out the group are the non-core product lines. Avride develops autonomous-driving technology and sidewalk delivery robots, with commercial partnerships including Uber and Uber Eats, Grubhub (at Ohio State University), Hyundai (for autonomous vehicles), and Rakuten, and testing across the U.S., South Korea, and Japan. TripleTen is an ed-tech platform that reskills and upskills individuals for technology careers through bootcamp and MOOC-style programs, operating in the U.S. and Latin America and named Best Software Bootcamp in the U.S. by Fortune. Toloka provides human-in-the-loop data services for AI training and evaluation, and ClickHouse offers a fast open-source real-time analytics database. While these do not drive the core investment thesis, they broaden the group’s exposure to AI-adjacent value and, critically, serve as monetizable funding sources for the capital-intensive cloud.

Figure 8 — Product & Platform Portfolio

Source: Company filings; Task 2/3 model & valuation. Prepared Jul 27, 2026.

Figure 8 Product & platform portfolio. Source: company disclosures.

Customers & Go-to-Market


Nebius serves three broad customer segments. At the top are hyperscalers and large platform companies contracting dedicated, multi-year capacity Microsoft and Meta are the anchor names, and their combined contracts exceed \$40 billion in backlog. These deals are investment-grade, long-duration, and prepayment-supported, providing revenue visibility and helping finance the build-out, but they also concentrate a large share of contracted revenue in a handful of counterparties. The second segment is AI-native companies model developers, coding and creative-AI startups, and agent builders with an early roster including Mistral, Cursor, Black Forest Labs, HeyGen, and Potoroom. The third is digital enterprises adopting AI at scale, such as Shopify, Prosus, ServiceNow, Cloudflare, and others migrating AI workloads onto specialized infrastructure.

The go-to-market blends direct enterprise sales with a growing partner ecosystem. Large capacity deals are negotiated directly and structured deal-by-deal, with management emphasizing margin discipline and system-level economics signing anchor contracts specifically to help finance the surrounding build-out. For the broader developer and enterprise base, Nebius combines self-service cloud consumption with technical partnerships (for example, LangChain integrations that route agent workloads onto Nebius Token Factory) and channel distribution through TD SYNEX, which exposes dedicated GPU capacity to resellers and enterprises seeking to unlock stalled AI deployments. Industry events NVIDIA GTC, high-growth partner conferences, and Nebius's own developer programs are important demand-generation channels.

The dominant commercial dynamic is scarcity. Across recent quarters Nebius has consistently sold out available capacity, with demand exceeding supply, pricing strengthening, and customers committing to larger and longer contracts. This inverts the typical enterprise-software sales motion: rather than chasing demand, Nebius is racing to convert contracted electrical power and GPU allocations into deliverable, revenue-generating capacity fast enough to satisfy a waiting order book. The strategic implication is that Nebius's growth is currently gated by supply data-center construction, power availability, and GPU delivery schedules more than by sales execution, which reframes the key operational question from "can they sell it?" to "can they build and deliver it on time and on budget?"

Sustaining a diversified customer base beyond the two hyperscaler anchors, and successfully re-selling the discretionary tranches of the Meta agreement to third parties, will be important markers of durable, high-quality growth.

Figure 9 — Customer Concentration by Contracted Value



Source: Nebius contract disclosures (Microsoft, Meta); AI-native/enterprise estimated, Jul 27, 2026.

Figure 9 Customer concentration by contracted value. Source: contract disclosures; AI-native/enterprise estimated.

Industry Overview

Nebius competes in the market for AI compute infrastructure the data centers, GPUs, networking, and cloud software required to train and run modern AI models. Within this, it belongs to the “neocloud” segment: specialist GPU-as-a-service (GPUaaS) providers built entirely around high-performance accelerated computing for AI and high-performance-computing workloads, as distinct from the general-purpose hyperscalers. Neoclouds emerged as an architectural response to a structural mismatch: traditional hyperscale systems were designed for generalized elasticity, whereas AI training imposes rigid requirements around massive parallelism, low-latency interconnect, locality, and dense concentration of compute. Specialist providers offer faster access to the latest GPUs, often at lower prices and with higher utilization, filling a gap that general-purpose clouds have struggled to serve at the scale AI demands.

The market is expanding at an extraordinary pace. Industry analysts estimate neocloud revenue reached roughly \$23 billion in 2025 (more than tripling year over year) and project it toward approximately \$20–25 billion of run-rate in 2026, rising toward \$180 billion by 2030; one widely cited forecast (Synergy Research Group) sees the broader neocloud market approaching \$400 billion by 2031, while GPU-as-a-service is variously projected to grow at roughly a 36% compound annual rate over the coming years. This demand is underpinned by unprecedented capital expenditure from the largest technology companies: hyperscalers including Amazon, Alphabet, Microsoft, and Meta are collectively expected to spend on the order of \$600–700 billion on AI-related infrastructure in 2026, with Meta alone guiding to as much as \$115–135 billion. Crucially, these giants increasingly cannot build capacity fast enough themselves, driving them to commit tens of billions of dollars to neocloud partners Microsoft alone has committed more than \$60 billion across such partnerships.

Several structural features shape the industry’s economics. First, it is intensely capital-intensive: providers must fund enormous up-front investment in GPUs, buildings, and power, and depreciation on rapidly deployed GPU

fleets consumes roughly half of revenue at the leading players, so none of the major neoclouds is yet profitable on a GAAP basis. Second, access to two scarce inputs cutting-edge NVIDIA GPUs and electrical power is the binding constraint on growth; securing gigawatts of contracted power and preferential GPU allocations is a core competitive activity. Third, financing is both an enabler and a systemic risk: much of the build-out is funded through GPU-collateralized debt, and analysts have flagged a “refinancing wall” as tens of billions of dollars of such loans across CoreWeave, Nebius, Lambda, Crusoe, and Applied Digital come due within a narrow 2026–2028 window among a small, correlated group of lenders.

The industry’s technology cadence is set largely by NVIDIA’s product roadmap. The transition from Hopper (H100/H200) to Blackwell and Blackwell Ultra, and next to the Vera Rubin generation, both drives demand (each generation unlocks new capability and efficiency) and creates obsolescence risk (older fleets depreciate against superior new silicon). NVIDIA’s own strategy of investing directly in neoclouds including its \$2 billion stake in Nebius and stakes in peers has attracted scrutiny as a form of “circular financing,” in which the dominant GPU supplier invests in the customers that buy its chips. Regulatory and geopolitical factors also matter: export controls on advanced chips, energy and permitting constraints on data centers, and sustainability pressures around power and water use all bear on the pace and location of build-out. Finally, the competitive perimeter is fluid the hyperscalers are simultaneously the neoclouds’ largest customers and their most formidable potential competitors, a tension that crystallized in mid-2026 around reports of a potential “Meta Compute” offering.

These structural features have fueled an active debate about whether neoclouds are durable technology platforms or leveraged bets on power and real estate wrapped in an AI narrative. Skeptics point to GAAP losses, GPU-collateralized debt, and the circularity of NVIDIA investing in its own customers; optimists counter that the demand is real, contracted, and prepaid by the most creditworthy companies in the world, and that operators with genuine software differentiation and disciplined financing will compound for years. Nebius sits in the middle of this debate as a relatively clean, well-capitalized operator with a full-stack strategy but it is not immune to the sector-wide risks, and its equity trades with the high volatility (beta above 3) characteristic of the group. Where an investor lands on the broader debate largely determines how they view NBIS.

Figure 15 — Neocloud Market Size Evolution (TAM)




Figure 15 Neocloud / GPUaaS market size evolution (TAM). Source: Synergy Research, industry estimates.

Competitive Landscape


Nebius competes across two fronts: against fellow neoclouds for the specialist AI-compute market, and against the hyperscalers whose general-purpose clouds remain the default for most enterprise workloads. The table below summarizes the principal specialist competitors.

The most direct comparison is CoreWeave, the segment leader and the other principal listed pure-play. CoreWeave is substantially larger roughly ten times Nebius's GPU fleet, FY2025 revenue of about \$5.13 billion, and a revenue backlog near \$66.8 billion and occupies the top independent "Platinum" tier, with Nebius in the "Gold" tier alongside Oracle and Azure. The two companies are, however, pursuing different strategies. CoreWeave has leaned into raw scale, whereas Nebius is building a fuller software stack (Token Factory, agentic search, inference optimization) on top of its compute, betting that enterprises will value an integrated platform rather than commodity GPU rental. On several financial-quality measures Nebius screens favorably: a cleaner balance sheet with lower leverage, a lower price-to-book multiple, and the optionality of monetizable non-core stakes, in exchange for smaller current scale.

Nebius's principal competitive advantages are fourfold. First, engineering depth: an in-house team of more than a thousand infrastructure and AI engineers with a long track record of building world-class systems, enabling differentiated price-performance and rapid iteration up the stack. Second, the NVIDIA relationship: a \$2 billion equity investment and joint build-out plan signal preferential access to scarce GPUs and next-generation platforms such as Vera Rubin. Third, financial flexibility: hyperscaler prepayments, low-coupon converts, asset-backed facilities, and non-core stakes together fund growth with comparatively modest dilution and leverage. Fourth, a European domicile and multi-region footprint that appeals to customers with data-sovereignty requirements and helps diversify power and regulatory exposure.


The vulnerabilities are equally clear. Nebius is sub-scale relative to CoreWeave and to the hyperscalers, and the underlying GPU-rental function is at risk of commoditization, which is precisely why the company is racing to add differentiated software. It is heavily dependent on two anchor customers and on NVIDIA. And the entire neocloud cohort is exposed to correlated financing risk and to the possibility that its largest customers vertically integrate. The mid-2026 report that Meta was exploring its own "Meta Compute" cloud which briefly cut around 15% from Nebius and CoreWeave shares is the sharpest illustration of the structural tension: the biggest customer could become the biggest competitor. Nebius's answer is to move up the value chain faster than the compute commoditizes, and to diversify its book beyond the anchor contracts.

Figure 16 — Competitive Positioning (bubble = EV)



Source: Company filings, market data ~Jul 2026 (smaller peers estimated). Prepared Jul 27, 2026.

Figure 16 Competitive positioning (bubble = EV). Source: company filings, market data (~Jul 2026).

Figure 17 — Listed-Neocloud Revenue Share (2026E)

Source: 2026E revenue estimates; illustrative of listed pure-plays only. Prepared Jul 27, 2026.

Figure 17 Listed-neocloud revenue share, 2026E (illustrative of listed pure-plays). Source: 2026E estimates.

Market Opportunity (TAM)


The addressable opportunity for Nebius is anchored in the broader build-out of AI compute. At the widest framing, hyperscalers alone are expected to deploy roughly \$600–700 billion of AI-related capital expenditure in 2026, a figure that captures the sheer scale of infrastructure being constructed and the portion that flows to specialist providers when the giants cannot build fast enough. The more directly relevant total addressable market is the neocloud / GPU-as-a-service segment: industry estimates put neocloud revenue at roughly \$20–25 billion in 2026, scaling toward approximately \$180 billion by 2030, with Synergy Research Group projecting the market to approach \$400 billion by 2031 and other analysts modeling GPU-as-a-service growth at roughly a 36% compound annual rate. Even conservative readings imply a market expanding several-fold over the current decade.

Against that TAM, Nebius's serviceable opportunity is defined by how much contracted power and GPU capacity it can bring online. The company exited 2025 with 170 MW of active power and roughly \$1.25 billion of ARR, has contracted more than 4 GW of power, and targets 800 MW–1 GW of connected capacity by the end of 2026 en route to a longer-term ambition supported by the NVIDIA partnership of as much as 5 GW by 2030. Management's

own near-term markers translate this capacity into financial targets of \$3.0–3.4 billion of 2026 revenue and \$7–9 billion of exit-2026 ARR, more than half of which was already contracted at the time of guidance. If the company delivers on its 2030 power ambition at anything like current asset productivity (roughly \$0.9 million of ARR per active megawatt on recent disclosures), the implied revenue base would be a multiple of today's.

The key variables governing how much of this opportunity Nebius actually captures are execution and supply, not demand. Demand is, for now, effectively unlimited relative to available capacity the company sells everything it builds. The binding constraints are the pace of data-center construction, the availability of electrical power, GPU delivery schedules, and the capital to fund a build-out that management now guides toward \$20–25 billion in 2026 alone (with cumulative multi-year spending running well into the tens of billions). Realistic share capture therefore depends on Nebius maintaining preferential GPU access, securing power at scale, financing the build efficiently, and converting contracted capacity into active, revenue-generating capacity on schedule. Success would make Nebius one of a small number of at-scale independent AI-cloud platforms; shortfalls in delivery, financing, or customer diversification would compress the realizable share of an otherwise vast market.

Figure 31 — Peer EV/Revenue Multiples



Source: Company filings, market data ~Jul 2026 (smaller peers estimated). Prepared Jul 27, 2026.

Figure 31 Peer EV/Revenue multiples. Source: company filings, market data (~Jul 2026).

Financial Analysis

Company snapshot


Metric	Value / description
Ticker / exchange	NBIS NASDAQ (Class A); ISIN NL0009805522
Headquarters	Amsterdam, Netherlands
FY2025 revenue	\$529.8M (+479% YoY, restated); Q4 group adj. EBITDA turned positive
Q1 2026 revenue	\$399.0M (+684% YoY, +75% QoQ); Nebius AI adj. EBITDA margin ~45%
Core AI-cloud ARR	\$1.25B exit-2025 → \$1.92B at Q1 2026 (+54% QoQ)
FY2026 guidance	Revenue \$3.0–3.4B; ARR \$7–9B; ~40% adj. EBITDA margin; capex \$20–25B
Contracted power	>4 GW contracted; 170 MW active YE2025 → 800 MW–1 GW target YE2026
Key contracts	Microsoft up to \$19.4B (through 2031); Meta up to \$27B (from 2027)
Strategic investor	NVIDIA \$2B equity (~9.3% stake); JV toward ~5 GW by 2030

Source: Nebius Q4 2025 & Q1 2026 results; SEC filings. As of Jul 27, 2026.

Historical financial analysis

Nebius's reported financials reflect a company barely two years into its life as an independent entity. Following the 2024 divestiture of the Russian Yandex operations, the continuing business generated revenue of \$91.5M in FY2024A (restated to exclude Toloka, now discontinued operations), which surged 479% to \$529.8M in FY2025A. Growth accelerated through the year: quarterly revenue climbed from \$55.3M in Q1 2025 to \$227.7M in Q4 2025 (+547% YoY), before reaching \$399.0M in Q1 2026 (+684% YoY). The core AI-cloud business drove this trajectory, exiting 2025 with \$1.25B of ARR comfortably ahead of the \$0.9–1.1B management had guided and reaching \$1.92B by Q1 2026.


Gross margin (which Nebius presents excluding depreciation) reached 68.6% in FY2025, reflecting the favorable economics of selling scarce GPU capacity under multi-year contracts. Below the gross-profit line, however, the picture is dominated by the capital intensity of the build-out. Depreciation and amortization rose from \$29.3M in FY2023 to \$77.1M in FY2024 and \$417.9M in FY2025 a 442% increase in the final year as the GPU fleet and data-center footprint expanded. As a result, FY2025 group operating margin was approximately –115.5% (an operating loss near –\$612M), even as the group reached breakeven adjusted EBITDA on an exit basis.

Figure 10 — Gross Margin Evolution (excl. D&A)

Source: Company filings; Task 2/3 model & valuation. Prepared Jul 27, 2026.

Figure 10 Gross margin evolution (excl. D&A), 2024A–2030E.

Reported GAAP net income from continuing operations was a modest positive \$29.0M in FY2025 (\$101.7M including discontinued operations) despite a \$596.2M operating loss a quirk driven by a \$598.9M non-cash gain on the company's equity stakes, principally the sharp re-rating of ClickHouse. This distinction matters for how the business should be read: reported net income is materially influenced by fair-value swings on non-core holdings (a \$621M gain in Q1 2026 vs a \$250M loss in Q4 2025), so adjusted EBITDA and ARR are the cleaner performance metrics. Cash generation in FY2025 benefited from large customer prepayments Q1 2026 operating cash flow reached \$2.3B but capital expenditure of \$4.07B in FY2025 (Q4 alone \$2.06B) drove free cash flow to -\$3.66B. The company exited 2025 with \$3.68B of cash and \$4.13B of convertible debt (a modest net-debt position), and by Q1 2026 held ~\$9.3B of liquidity after the NVIDIA raise.

Figure 12 — Free Cash Flow Trend (CFO – Capex)

Source: Company filings; Task 2/3 model & valuation. Prepared Jul 27, 2026.

Figure 12 Free cash flow trend (CFO – capex), 2024A–2030E. Source: Task 2 model.

Income statement (\$M)	FY2024A	FY2025A	FY2026E	FY2027E	FY2028E	FY2029E	FY2030E
Revenue	92	530	3,200	7,500	12,000	16,500	21,000
Cost of revenue (excl. D&A)	(44)	(166)	(1,024)	(2,325)	(3,600)	(4,950)	(6,300)
Gross profit	48	363	2,176	5,175	8,400	11,550	14,700
Product development	(115)	(177)	(384)	(825)	(1,200)	(1,568)	(1,890)
Sales, general & administrative	(256)	(379)	(608)	(1,260)	(1,872)	(2,310)	(2,814)
D&A	(77)	(404)	(1,760)	(3,750)	(5,520)	(7,095)	(8,400)
Loss from operations (EBIT)	(400)	(596)	(576)	(660)	(192)	578	1,596
Memo: adjusted EBITDA	(226)	(65)	1,280	3,300	5,640	8,085	10,500
Net income (continuing)	(352)	29	(612)	(1,397)	(1,506)	(1,041)	(132)

Source: reported filings (FY2024A restated ex-Toloka; FY2025A audited) and rebuilt model for projections. Figures in \$M.

Projection Assumptions


Our projections cover FY2026E–FY2030E and are anchored on management guidance where available and on a capacity-driven build-out thereafter. Given Nebius’s limited operating history, we place particular weight on the two hard data points management has provided for FY2026 revenue of \$3.0–3.4B and exit ARR of \$7–9B and then extend the trajectory using contracted power (>4 GW, scaling toward ~5 GW by 2030) and the >\$40B contracted backlog. The single most important modeling judgment is that demand is not the constraint; the pace at which

contracted capacity converts into revenue-generating capacity is. Our base case assumes Nebius executes broadly in line with its stated ambitions without either the acceleration of the bull case or the slippage of the bear case.

A. Revenue by segment

Core AI cloud. We project core AI-cloud revenue to grow from \$450M in FY2025A to roughly \$20.4B in FY2030E, a ~112% five-year CAGR, lifting the core from ~85% of group revenue in FY2025 to ~97% by FY2030. The build is anchored on exit ARR, which we ramp from \$1.25B (2025A) to \$8.0B (2026E, the midpoint of \$7–9B guidance) and then to \$13.5B, \$18.0B, \$22.0B and \$26.0B across 2027E–2030E. The specific drivers are: (1) conversion of the Microsoft (up to \$19.4B) and Meta (up to ~\$27B) contracts into delivered capacity, with Meta’s Vera Rubin-based tranches beginning in early 2027; (2) contracted power scaling from 170 MW active at YE2025 toward 800 MW–1 GW connected by YE2026 and ~5 GW by 2030; (3) firm-to-rising pricing as demand persistently exceeds supply; and (4) a growing on-demand book from AI-native companies and digital enterprises. By year, we assume FY2026E group revenue of \$3.2B (guidance midpoint), then \$7.5B (2027E), \$12.0B (2028E), \$16.5B (2029E) and \$21.0B (2030E) a deceleration in percentage terms (504% → 27%) that still compounds at ~110% over five years. The chief risk to these assumptions is delivery: gigawatt-scale construction and GPU timing must stay on schedule, and the discretionary Meta tranches must be re-sold to third parties if not taken.

Figure 3 — Revenue by Segment (2024A-2030E) (KEY CHART)



Source: Company filings; Task 2/3 model & valuation. Prepared Jul 27, 2026.

Figure 3 Revenue by segment, 2024A–2030E (core AI cloud vs non-core).

Non-core (TripleTen, Avride). We project non-core revenue to grow modestly from ~\$80M in FY2025A to ~\$630M by FY2030E. TripleTen (ed-tech reskilling) and Avride (autonomous mobility) are real businesses with commercial partnerships Avride works with Uber, Grubhub, Hyundai and Rakuten but they are not the investment case and we assume they shrink as a share of the group from ~15% (2025) to ~3% (2030). We deliberately assign no explicit revenue to the ClickHouse and Toloka stakes, which are equity holdings rather than consolidated operating


segments; their value accrues through the balance sheet (and, historically, through large non-cash gains) rather than the income statement, and we treat them as a source of non-dilutive funding optionality.

Segment (\$M)	FY2024A	FY2025A	FY2026E	FY2027E	FY2028E	FY2029E	FY2030E
Core AI cloud	50	450	3,040	7,200	11,520	15,840	20,370
Core mix (% of group)	55%	85%	95%	96%	96%	96%	97%
Non-core (TripleTen, Avride)	41	79	160	300	480	660	630
Total group revenue	92	530	3,200	7,500	12,000	16,500	21,000
Core AI-cloud exit ARR	—	1,250	8,000	13,500	18,000	22,000	26,000

Core mix rises toward 97% as the AI-cloud business scales.

B. Revenue by geography

Nebius does not disclose a clean geographic revenue split, so the following mix is estimated and calibrated to the company's footprint and contract locations. We assume the revenue base shifts steadily toward North America as the large U.S. data-center commitments underpinning the Microsoft (Vineland, New Jersey) and Meta agreements come online. Europe. We project Europe's share to decline from ~55% of revenue in FY2025A to ~40% by FY2030E even as absolute European revenue grows from ~\$291M to ~\$8.4B, reflecting the company's Amsterdam and Finland roots, data-sovereignty demand, and continued build-out in France and the Nordics. North America. We project North America to rise from ~35% to ~52% of revenue over the same period (from ~\$185M to ~\$10.9B), driven by the hyperscaler contracts anchored on U.S. facilities, the new 1.2 GW Pennsylvania AI factory, and the Bay Area engineering presence added through the Eigen AI acquisition. Middle East & Rest of World. We hold this at ~8–10%, reflecting selective expansion (including the Middle East) without assuming major new regional commitments. By 2030E, North America becomes the largest region a meaningful shift from the company's European origins and a direct consequence of where the anchor contracts are being delivered.

Figure 4 — Revenue by Geography, estimated mix (2024A-2030E) (KEY CHART)

Source: Geographic mix estimated (company does not disclose a clean split); Task 2 model. Prepared Jul 27, 2026.

Figure 4 Revenue by geography (estimated mix), 2024A–2030E

Geography	FY2025A	FY2026E	FY2027E	FY2028E	FY2029E	FY2030E
Europe (% of revenue)	55%	50%	46%	43%	41%	40%
North America (%)	35%	42%	46%	49%	51%	52%
Middle East & RoW (%)	10%	8%	8%	8%	8%	8%
Europe revenue (\$M)	291	1,600	3,450	5,160	6,765	8,400
North America revenue (\$M)	185	1,344	3,450	5,880	8,415	10,920
Middle East & RoW (\$M)	53	256	600	960	1,320	1,680

Geographic mix is estimated; company does not disclose a clean split.

C. Other key assumptions


Gross margin. We hold gross margin (excluding D&A) at ~68–70% across the forecast, roughly in line with the 68.6% reported in FY2025 and consistent with the economics of selling scarce, contracted GPU capacity.

Operating expenses. We model product development (R&D) declining from ~13% of revenue in FY2026E to ~9% by FY2030E, and sales, general & administrative expense from ~15% to ~11%, as fixed costs are spread over a rapidly growing revenue base. Calibrated together with gross margin, these produce a group adjusted EBITDA margin that rises from ~40% (matching FY2026 guidance) to ~50% by FY2030E.

Depreciation & amortization. This is the pivotal assumption. We model D&A as a declining percentage of revenue (55% → 40% across the forecast), calibrated to the placed-in-service asset base construction-in-progress is not depreciated until capacity goes live, so reported D&A lags the enormous capital expenditure. Even so, D&A consumes roughly half of revenue throughout, which is why GAAP EBIT stays negative until 2029E. Capital

expenditure. We assume FY2026E capex of \$22.5B (the midpoint of \$20–25B guidance), moderating to \$18.0B, \$12.0B, \$9.0B and \$8.0B across 2027E–2030E as the build matures a cumulative ~\$69.5B over five years. Working capital. Customer prepayments (deferred revenue) are a major source of operating cash; we grow the deferred-revenue balance in line with new mega-contract prepayments before it plateaus as capacity is delivered and recognized. We assume DSO of ~45 days and DPO of ~60 days. Financing and interest. The build is funded by prepayments, convertible notes, asset-backed debt and the \$2B NVIDIA equity investment; we compute interest on opening balances to avoid circularity and apply a blended ~4.0% rate. Tax. We apply a 15% effective rate on positive pre-tax income (below the ~25.8% Dutch statutory rate, reflecting NOL carryforwards) and zero when pre-tax income is negative. Equity-stake gains. We conservatively model zero gains from ClickHouse and other holdings in the forecast, even though such gains turned FY2025 GAAP net income positive representing lumpy upside not captured in our base case.

Figure 23 — Capex vs Depreciation & Amortization



Source: Company filings; Task 2/3 model & valuation. Prepared Jul 27, 2026.

Figure 23 Capex vs depreciation & amortization, 2024A–2030E.

Scenario Analysis

We frame the investment case around three scenarios. The wide spread of outcomes reflected in a DCF range of roughly \$129–\$291 and a comps range of \$144–\$233 is itself a central feature of the Nebius story: the equity is valued almost entirely on cash flows beyond the explicit forecast, so relatively modest changes in the growth path, terminal margin and exit multiple swing fair value substantially.

Bull Case asset-light FCF inflection and share gains (probability ~25%)


Key assumptions: 2025–2030E revenue CAGR of ~125% (vs ~109% base), lifting FY2030E revenue to ~\$30.6B; a FY2030E adjusted EBITDA margin of ~54% (vs ~50% base), implying ~\$16.5B of adjusted EBITDA; faster power and GPU delivery; and, critically, a successful asset-light pivot that shifts facility financing and ownership to infrastructure partners, materially reducing on-balance-sheet capital intensity and lifting free cash flow well above our base case. Catalysts required: (1) additional multi-billion-dollar hyperscaler or frontier-lab contracts beyond Microsoft and Meta (timing: 2026–2027); (2) tangible evidence that the asset-light model improves unit FCF (2026–2027 reporting); (3) accelerated ARR beats versus the \$7–9B FY2026 exit guidance; and (4) non-dilutive capital

from ClickHouse or Toloka monetization. Rationale: in this scenario Nebius converts its backlog faster and at better economics than modeled, re-sells the discretionary Meta tranches to third parties at strong pricing, and the market rewards the improved cash-flow profile with a higher terminal multiple. The full-stack software layer (Token Factory, agentic search) gains traction, differentiating Nebius from commoditizing GPU rental and supporting durable premium pricing. Valuation implications: at a ~17x terminal exit multiple and an ~11% WACC, our DCF supports roughly \$255–\$291 per share; growth-justified comps (~9x EV/2027E revenue) point to ~\$233. We set a bull-case value of approximately \$260 per share (~31% upside).

Base Case guidance delivered (probability ~55%)

Key assumptions: 2025–2030E revenue CAGR of ~109%, reaching ~\$21.1B in FY2030E; a FY2030E adjusted EBITDA margin of ~50% (~\$10.6B of adjusted EBITDA); FY2026 revenue and ARR delivered in line with the \$3.0–3.4B and \$7–9B guidance ranges; capex of ~\$69.5B cumulative over 2026–2030E; and GAAP profitability remaining elusive through 2030E on heavy D&A and interest, with adjusted EBITDA strongly positive. Rationale: this is simply management executing its stated plan without material acceleration or slippage. Demand remains ample, capacity converts roughly on schedule, and the two anchor customers honor their commitments. Valuation: our DCF base case is \$205 per share (WACC 11.4%, 15x terminal exit multiple), and growth-justified trading comps (~7.25x EV/2027E revenue) yield \$189. Weighting DCF 60% and comps 40% produces a base-case 12-month target of \$198 per share essentially the current price, which is the core of our Neutral rating.

Figure 14 — Scenario Comparison (2030E)



Source: Company filings; Task 2/3 model & valuation. Prepared Jul 27, 2026.


Figure 14 Scenario comparison, 2030E revenue and adjusted EBITDA.

Bear Case in-sourcing, delivery slippage and financing stress (probability ~20%)

Key assumptions: 2025–2030E revenue CAGR of ~85% (vs ~109% base), holding FY2030E revenue to ~\$11.5B; a FY2030E adjusted EBITDA margin of ~42% (vs ~50%), implying ~\$4.8B of adjusted EBITDA; slower capacity ramp; and pressured pricing. Downside triggers: (1) Meta and/or Microsoft slow, delay or renegotiate the mid-2026 “Meta Compute” scare, which cut ~15% from the sector, showed how quickly a large customer becoming a competitor can de-rate neoclouds (likelihood: moderate); (2) gigawatt-scale delivery slips against contractual dates, triggering penalties or termination rights (moderate); (3) the 2026–2028 refinancing wall across the neocloud cohort coincides with tighter credit, raising financing costs or constraining the build (moderate); and (4) GPU-rental

commoditization compresses margins before the software stack gains scale (lower). Rationale: in this scenario Nebius's growth still looks impressive in absolute terms, but the market re-rates the shares sharply because the valuation had priced in flawless execution. With most DCF value in the terminal, a lower exit multiple (~12x) and higher WACC (~12.5%) compress fair value quickly. Valuation implications: our DCF supports roughly \$110–\$149 per share under these parameters, and comps (~5.5x EV/2027E revenue) point to ~\$144. We set a bear-case value of approximately \$140 per share (~29% downside).

Figure 33 — Price Target by Scenario (vs current)



Source: Company filings; Task 2/3 model & valuation. Prepared Jul 27, 2026.

Figure 33 Price target by scenario vs current price.


Scenario comparison and risk/reward. On a probability-weighted basis (Bull 25% at ~\$260, Base 55% at ~\$198, Bear 20% at ~\$140), the expected value is approximately \$200 per share consistent with our \$198 target and Neutral rating. The distribution is roughly symmetric around the current price, with ~31% upside in the bull case against ~29% downside in the bear case. Because the outcome is highly path-dependent hinging on delivery timing, customer behavior and financing conditions we see balanced risk/reward at today's price and would require either a lower entry point or clearer evidence of the asset-light FCF inflection to move constructive.

Scenario	Probability	2030E revenue	2030E adj. EBITDA	EBITDA margin	Value/share
Bull	~25%	\$30.6B	\$16.5B	54%	~\$260
Base	~55%	\$21.1B	\$10.6B	50%	\$198
Bear	~20%	\$11.5B	\$4.8B	42%	~\$140
Probability-weighted	100%	—	—	—	~\$206

Growth Drivers

We see five principal drivers of the growth trajectory embedded in our model. First, hyperscaler backlog conversion: the Microsoft (up to \$19.4B) and Meta (up to ~\$27B) contracts alone underpin more than \$40B of contracted revenue, and converting this backlog into delivered, revenue-generating capacity is the single largest driver most of the FY2027E–FY2030E revenue ramp is contracted rather than speculative. Second, power and GPU scaling: Nebius has contracted more than 4 GW of power and targets ~5 GW by 2030 under its NVIDIA partnership; at anything like current asset productivity, each incremental gigawatt supports a large increment of ARR, making power procurement and construction pace the binding growth levers. Third, the software and inference stack: Token Factory, Tavily and Eigen AI move Nebius up the value chain toward higher-margin managed inference and agentic workloads, which both improves unit economics (more tokens per GPU) and differentiates the offering from commodity rental. Fourth, customer diversification: broadening beyond the two anchor customers and re-selling the discretionary Meta tranches to third parties would raise the quality and durability of the revenue base. Fifth, non-core monetization: the ClickHouse and Toloka stakes provide a source of non-dilutive capital (and, historically, large non-cash gains) that can fund the capital-hungry core without additional equity issuance. Each driver is quantified in the model; together they support the ~110% five-year revenue CAGR in our base case, while also defining the areas where slippage would most damage the thesis.

Figure 26 — Funding Sources for the Build-out



Source: Nebius disclosures (converts, NVIDIA equity, secured facility, Q1 2026 liquidity); prepayments estimated. Jul 27, 2026.

Figure 26 Funding sources for the build-out. Source: company disclosures; prepayments estimated.


Figure 13 — Operating Metrics Dashboard

Figure 13 Operating metrics dashboard. Source: Task 2 model.

Valuation

We value Nebius using a discounted cash flow analysis (60% weight) and trading comparables (40% weight), with precedent transactions used only for control-premium context. The methods converge on a 12-month price target of \$198 per share, right on top of the current price the basis for our Neutral rating. The sections below reproduce the core of our valuation work; full detail sits in the accompanying Excel model.

Discounted cash flow

We value Nebius on unlevered free cash flow discounted at an 11.4% weighted average cost of capital. The WACC reflects a CAPM cost of equity of 13.1% (4.3% risk-free rate, 1.6 levered beta, 5.5% equity risk premium), an after-tax cost of debt of 4.7%, and an 80/20 equity/debt weighting a level consistent with a high-growth, capital-intensive infrastructure business. Free cash flow is drawn directly from the model: unlevered FCF of -\$18.4bn (2026E), -\$13.1bn (2027E), -\$6.0bn (2028E), -\$1.2bn (2029E) and +\$1.4bn (2030E), reflecting a build-out in which capital expenditure runs far ahead of revenue before the asset base begins to monetize.

Because year-five free cash flow is not yet normalized, we use an exit-multiple terminal value (15x 2030E adjusted EBITDA of ~\$10.5bn) as our primary method; the 15x multiple sits below CoreWeave's current ~18x, reflecting a more mature, slower-growth business by 2030. A perpetuity-growth cross-check is uninformative here with

depressed 2030E FCF it implies a negative value which itself signals that the explicit forecast window is shorter than the company's investment cycle. The table summarizes the base case.

Sensitivity

Given the terminal dependence, we stress the two variables that matter most WACC and the terminal exit multiple. Across a WACC range of 9.9%–12.9% and exit multiples of 11x–19x, the implied value per share spans roughly \$110 to \$313, with the base case at \$205. A ± 100 bps move in WACC shifts value by roughly \$15–20 per share; a ± 2 x move in the exit multiple shifts it by roughly \$40 per share. The heat-mapped tables are provided in the accompanying Excel model (Sensitivity tab).

DCF base case	Value
Sum of PV of explicit uFCF (2026–30E)	(\$32,628)M
Terminal value (15x 2030E adj. EBITDA)	\$157,500M
PV of terminal value	\$91,803M
Enterprise value	\$59,174M
Terminal value as % of EV	155%
Add: net cash	\$1,500M
Equity value	\$60,674M
Diluted shares	296M
Implied value per share	\$205

WACC 11.4%; terminal exit multiple 15x 2030E adjusted EBITDA.




Figure 28 DCF sensitivity: value/share by WACC and terminal exit multiple.

WACC \ Exit x	11x	13x	15x	17x	19x
9.9%	\$142	\$186	\$230	\$274	\$319
10.9%	\$132	\$175	\$217	\$259	\$302
11.4% (base)	\$122	\$164	\$205	\$246	\$288
11.9%	\$124	\$164	\$205	\$245	\$285
12.9%	\$116	\$154	\$193	\$232	\$270

Base case \$205 (WACC 11.4%, 15x).

Comparable companies

The listed neocloud peer set is small and, apart from CoreWeave, dominated by miner-to-AI pivots (IREN, Applied Digital, Core Scientific, TeraWulf) whose economics are only partly comparable. Because none of these companies is meaningfully profitable on a GAAP basis, EV/Revenue is the appropriate multiple; we anchor on forward (2027E) revenue, where the extreme near-term growth differentials partially normalize. Market data are as of late July 2026; figures for the smaller peers are estimated.

The comps reveal the crux of the debate. On EV/2027E revenue, Nebius trades at 7.6x versus a peer median of ~4.2x a ~2x premium. On unadjusted peer medians, NBIS would be worth only ~\$100–120 per share, implying the stock is expensive. We do not think a pure peer-median is the right lens: Nebius grows several times faster than the group, carries a cleaner balance sheet (strong liquidity, prepayment-funded build, and a recently announced asset-light pivot that shifts facility financing to partners), and differentiates through a full-stack software layer (Token Factory, agentic search, managed inference). Applying a growth-justified 7.25x EV/2027E revenue roughly where Nebius trades today and ~1.9x the peer median yields a base comps value of \$189 per share, within a \$144 (5.5x) to \$233 (9.0x) range. Whether that premium is warranted is precisely what separates the bulls from the bears.

Company	EV (\$bn)	Rev 2027E	EV/26E	EV/27E	Rev growth
CoreWeave (CRWV)	72.9	\$19.0bn	5.8x	3.8x	112%
IREN (IREN)*	8.5	\$2.5bn	7.1x	3.4x	~150%
Applied Digital (APLD)*	7.5	\$1.6bn	10.7x	4.7x	~120%
Core Scientific (CORZ)*	6.8	\$1.6bn	7.6x	4.3x	~60%
TeraWulf (WULF)*	4.6	\$1.1bn	9.2x	4.2x	~100%
Nebius (NBIS)	56.7	\$7.5bn	17.7x	7.6x	684%
Maximum (peers)	—	—	10.7x	4.7x	—
75th percentile	—	—	—	4.5x	—
Median	—	—	7.6x	4.2x	—
25th percentile	—	—	—	3.8x	—
Minimum (peers)	—	—	5.8x	3.4x	—

*Estimated. Market data ~Jul 24–27, 2026. Statistical summary computed on peer EV/2027E revenue (NBIS excluded).

Trading comps (EV/2027E Rev)	\$144	\$189	\$233	40%
Precedent transactions	context	context	context	0%
Weighted 12-month price target		\$198		100%

Source: Task 3 valuation. Precedents carry no formal weight (M&A of a ~\$59bn listed neocloud is unlikely near-term).

We rate Nebius NEUTRAL (Hold) with a \$198 12-month target. This is not a negative view of the business Nebius is among the highest-quality, best-financed independent AI-infrastructure platforms, with genuine full-stack differentiation and a powerful NVIDIA relationship. It is a statement about price: after a 262% run, the base case is already in the stock, and the payoff is skewed by an unusually wide, terminal-dependent value distribution. We would move to Buy on a pullback toward the \$150s, or on hard evidence that the asset-light pivot structurally lifts free cash flow (which would raise our DCF materially), and to a more cautious stance if customer-concentration or financing risks crystallize.

Key catalysts (upside)


- **Q2 2026 results (July 28, 2026).** A beat-and-raise on revenue/ARR or another large contract could re-rate the shares; the print lands the day after this note.
- **Asset-light FCF inflection.** If facility financing and ownership shift to infrastructure partners as announced, capital intensity falls and free cash flow improves a direct uplift to our DCF, which currently assumes a heavy on-balance-sheet build.
- **New hyperscaler / frontier-lab contracts.** Additional multi-billion-dollar capacity deals (beyond Microsoft and Meta) would extend the backlog and diversify concentration.
- **Non-core monetization.** Further value crystallization in ClickHouse, Toloka or Avride provides non-dilutive capital and optionality not captured in our base case.
- **Index / ownership tailwinds.** Recent Nasdaq-100 inclusion and NVIDIA's ~9.3% stake broaden the institutional ownership base.

Key risks (downside)

- **Terminal-value dependence.** With >100% of DCF value in the terminal, small changes in the exit multiple or WACC swing fair value by tens of dollars per share.
- **Customer concentration & Meta in-sourcing.** Microsoft and Meta dominate the backlog; the mid-2026 "Meta Compute" scare (which cut ~15% from the group) shows how quickly a large customer becoming a competitor can de-rate the sector.
- **Financing & the refinancing wall.** The build depends on prepayments, convertibles and asset-backed debt; a 2026–2028 refinancing wall across the neocloud cohort is a systemic risk if credit tightens.
- **Execution & delivery.** Contracts carry delivery-date provisions; slippage on gigawatt-scale construction or GPU timing could trigger penalties or termination rights.
- **Dilution & GAAP losses.** Share count has risen ~11% year over year, and on our conservative base case (excluding equity-stake gains) the company remains GAAP-lossmaking through 2030 on heavy D&A and interest.

Data as of July 27, 2026 (market close prior day; NBIS \$198.25). Valuation derived from the accompanying Task 2/Task 3 Excel model. Sources: Nebius FY2025 20-F and Q1 2026 results; StockAnalysis, Trefis, Yahoo Finance, Robinhood, Simply Wall St, CNBC (market data). This document is an analytical exercise and not investment advice; the author is not a licensed investment adviser.


Figure 34 — NBIS Forward EV/Revenue Multiple Compression



Source: EV/NTM-revenue, approximate from market data; multiple compresses as revenue scales. Jul 27, 2026.

Figure 34 NBIS forward EV/revenue multiple compression. Source: market data (approximate).


Figure 35 — Price Target Landscape



Source: TipRanks, Simply Wall St, Baird (~Jul 2026); range 120-410. Our target from Task 3. Jul 27, 2026.

Figure 35 Price target landscape (analyst context). Source: TipRanks, Simply Wall St, Baird (~Jul 2026);

Figure 27 — Ownership Structure (approx.)



Source: Company filings, TipRanks (Volozh ~11.5%, NVIDIA ~9.3%); institutions/float estimated. Jul 27, 2026.

Figure 27 Ownership structure (approximate). Source: company filings, TipRanks; institutions/float estimated.

Appendix A Financial Statements

Cash flow highlights

Cash flow (\$M)	FY2024A	FY2025A	FY2026E	FY2027E	FY2028E	FY2029E	FY2030E
Cash from operations	(270)	402	4,250	3,747	4,481	6,134	7,939
Capital expenditures	(808)	(4,066)	(22,500)	(18,000)	(12,000)	(9,000)	(8,000)
Free cash flow	(1,077)	(3,664)	(18,250)	(14,253)	(7,519)	(2,867)	(61)
Debt issuance, net	—	4,073	18,000	17,000	9,000	2,500	—
Equity issuance, net	657	1,052	2,000	500	300	200	200
Ending cash	2,435	3,678	4,228	7,175	8,756	8,489	8,528

Balance sheet highlights

Balance sheet (\$M)	FY2024A	FY2025A	FY2026E	FY2027E	FY2028E	FY2029E	FY2030E
Cash & equivalents	2,435	3,678	4,228	7,175	8,756	8,489	8,528
PP&E, net	847	5,573	26,313	40,563	47,043	48,948	48,548
Total assets	3,549	12,450	34,614	52,641	61,457	63,750	64,044
Total debt	6	4,128	22,128	39,128	48,128	50,628	50,628
Total equity	3,254	4,613	6,097	5,410	4,516	4,087	4,659

Appendix B Data Sources & References

Market and financial data as of the most recent disclosures available July 27, 2026 (Q1 2026 results reported May 13, 2026; NBIS close \$198.25 on Jul 27). Selected sources below; all figures are attributed in the underlying model and prior task documents.

- **Nebius investor relations & results:** Q4 2025 and Q1 2026 results, shareholder letters, and SEC filings (Microsoft, Meta, financing). <https://group.nebius.com/investors>
- **Nebius newsroom:** Meta agreement, Eigen AI and Tavily acquisitions, senior secured facility. <https://nebius.com/newsroom>
- **CoreWeave data (comps anchor):** Market cap, EV, revenue, and multiples (~Jul 2026). <https://stockanalysis.com/stocks/crwv/statistics/>
- **Nebius market data:** Price, market cap, shares, 52-week range (~Jul 2026). <https://stockanalysis.com/stocks/nbis/>
- **CoreWeave forward revenue / neocloud market:** 2026E revenue and EV/forward-revenue context. <https://www.trefis.com/stock/crwv/>
- **Analyst context:** Consensus targets, Baird initiation, NVIDIA stake coverage (TipRanks via CNN). <https://www.cnn.com/markets/stocks/NBIS>
- **Neocloud landscape & Meta in-sourcing:** Sector derating, competitive positioning, and market data. <https://finance.yahoo.com/quote/NBIS/>

Note: geographic mix, smaller-peer figures, ownership/float, prepayments, and the stock-price series are estimated or approximate where not separately disclosed, and are labeled as such throughout. This report is an analytical exercise and not investment advice;