

NEBIUS

**March 4, 2026
7:45 a.m. PST**

Josh Baer: All right, before we get started, for important disclosures, please see the Morgan Stanley research disclosure website at www.morganstanley.com/researchdisclosures. And if you have any questions, please reach out to your sales representative.

My name is Josh Baer, software analyst here at Morgan Stanley. We are thrilled to have the Nebius leadership team here today or part of it, Arkady Volozh and Marc Boroditsky. Thank you so much for joining us- CRO.

Arkady Volozh: Thank you.

Marc Boroditsky: Thank you.

Josh Baer: Maybe to start it off, Arkady, for those newer to the Nebius story, could you walk us through some of the history and what you're trying to build?

Arkady Volozh: Well, that's not the first project of our team. We have some history. The team of engineers who started this company 18 months ago, the previous 25 years was running one of the largest IT companies in Europe. Infrastructure-wise, it was 200-megawatts of different data center equipment. We were the largest NVIDIA partner outside of U.S., China. So we had some experience.

And then a lot of geopolitical events happened, and several hundred engineers actually went to Europe, and we decided to start a new company. We left in 2022. And then several months later, this ChatGPT moment happened. And we realize that probably there is something for us to do here because we know the area very well. We were in machine learning, not even AI yet for many years. We know how to run such systems, and we were one of the big tech companies ourselves. And we realized that there is a space in this new era, there will be a lot of -- all the big tech companies, all the big tech systems will go there, but also there will be a lot of independent companies in the enterprise market, which will go into this area. They all will need experienced infrastructure, people who can provide them with experience in infrastructure and who actually can make them competitive to the big tech companies. And we thought that this is probably an area where we could be useful.

We started with the company just on summer '24. When we started, we had just one data center with us. That's all we had, 25 megawatts. We managed to sell our assets for several -- \$2.5 billion. It was a big -- much bigger company, but whatever we managed to sell, it was nice to start with. So we started with 25 megawatts, 18 months ago. Today, we're running 10x of that, almost. We publicly announced that we will be running another 5x on that. It will be almost 1 gigawatt of connected power this year, 800 to 1 gigawatt. And we also announced that we are building much more now. We have a pipeline already reserved for two gigawatts, and it will be probably close to three gigawatts this year.

But what we are building is not just gigawatts. We're building a real platform for AI developers, a real cloud, full stack cloud. We build the basic layers starting from the ground, land power shell, we build our design and build our own data centers. We fill them with our own design racks. So on the first layer, we are like companies, I don't know, Iron, Applied Digital others who build the data centers. We have it internally. Then we build the racks, it's like, let's say, Supermicro, Dell, they

build their racks, we build our own racks. Then we integrate everything and provide the basic layer, bare metal layer, it's called to big players and whoever wants it. And on top of it, we have a product, which is our own multi-tenant cloud. This is the same kind of cloud as the hyperscalers have. The big names, AWS, Azure, GCP. These are examples of multi-tenant cloud. We have this full stack. We recreated the full stack. We -- again, we didn't inherit any IP rights. We had to recreate it from scratch. And that's why this whole cloud is very modern, very fresh with no incumbency. And this is our new stack. This is a new cloud with a lot of services on top of it. It's not just multi-tenant cloud, which means basic cloud services, multi-tenancy billing, security, storage, managed databases, it's all there. But on top of it, there is a lot of services for -- specific for AI builders like our Token Factory, inferencing platform or agentic search services, which people who build AI need, or just recently, we announced human experts in the loop there. Human experts who actually make- perform tasks given to them by agents. It's agents running humans today.

So that's examples of services which are on our platform. And on top of that, we have a sales machine, of course, go-to-market, which is verticalized in geographies, verticalized in industries like pharmacy, robotics, retail, whatever. So this is a full stack from ground, to racks, to clouds, to services and the sales channel, et al. This is the full stack cloud. And we build it in like, two-dimension, this functionality and scale. And in scale, as I said, it's gigawatt scale. We are -- if you compare us to the big hyperscalers, we are probably the smallest today. But we are the same -- like our cloud today provides hundreds of thousands of GPUs. How many GPUs each of the hyperscalers provides? It's not millions, it's also hundreds of thousands, maybe more hundreds, but the same class. So we are in the class of -- by scale, hundreds of cloud, full-scale cloud with hundreds of thousands of GPUs in there. We'll go to millions eventually when the gigawatts start materializing. So it's full stack and full scale.

Josh Baer: That's a great overview and such an interesting story for a relatively new company, start with a lot of cash, hundreds of engineers with a longer history of expertise to move fast and build this full stack cloud.

Marc, I want to bring you into the conversation, if you could talk a little bit about some of the different customer segments that you're prioritizing, and really what gives you confidence in addressing different -- those different groups?

Marc Boroditsky: Certainly. And as you all should have heard, Arkady just laid out a very bold and broad vision in terms of the tech stack that we're delivering and the intent to actually meet the hyperscalers competitively. As a complement, we are also implementing a very broad distribution strategy. Today, we have the privilege of servicing thousands of developers that come to us to build on our platform as well as having contracted with hundreds of customers. It ranges from individual AI researchers all the way up to some of the biggest companies in the world like Microsoft that could spend billions of dollars on our platform.

You could say that because of our AI builder orientation, we have a right to win with AI start-ups. So this has been a big part of our success in the recent past. We do that, as Arkady mentioned, by meeting them where they are with the technical capabilities that the AI builder is looking for. That same readiness has actually won us opportunities with major software vendors that are embedding AI in their offerings. So companies like Shopify and Cloudflare. And then even more recently, we started to see opportunities with the enterprise. We believe in the fullness of time, enterprise will represent the majority of the TAM. Today, it's relatively narrow. So we have a very focused workload-specific set of initiatives where we're pursuing specific verticals, including health care, life sciences, physical AI, media and entertainment, retail and e-tail as well as financial services.

Josh Baer: Excellent. Arkady, you laid out a pretty ambitious plan as far as capacity. Can you talk a little bit more about that expansion? I think yesterday, you announced your approval for your first gigawatt site. What needs to occur to reach those goals of three gigawatts and beyond?

Arkady Volozh: Yes. As I said, we started with just one data center, which came -- which was built outside of our core market in Finland. That was our first. Then we started quickly leasing data centers because for -- to save time, you have no time to build, you need to lease. The first year, we were leasing everything we could, and we launched locations in France, in the U.K., in Iceland, in the U.S., of course, and in Israel, that was our first year. And today, we would build -- we mostly built ourselves. And this year, we are launching sites, which belong to us, which we built from scratch. And these are much bigger sites. This was the first example, Independence, Missouri. It's 800-megawatt, expandable to 1.2 gigawatt. Just got approval from the local communities. We have more sites coming like that. We build -- continue building in Europe. Majority of what we build now is in U.S., but we build not only in Europe and the U.S., we are now opening our offices in Middle East. I don't know if we can mention this today. But yes, that was outlined last week. Singapore and Korea and Japan there as well.

But what we build today is sites, we happen to build the largest -- currently the largest sites in Europe. For us, it's not so big, it's 240 megawatt in France. But for today, it's the largest in Europe. In U.S., we are building much more, it's hundreds of megawatts, going to gigawatts side. This is the type of sites we are looking for. And we have dozens of them. Today, we have 16 something locations announced. Again, it's a pipeline. It's like 10, 20 is already announced, and it's in, I don't know, a dozen of different markets. It will be more. It will be a network of data centers. Having your own platform allows you to use your own big sites, but also integrate anything you can find. If you found 20, 40 megawatts data center somewhere, just plug it in, it works immediately get the software platform on top of it. All the sales machine works for it. And for the customers, it's just a cloud. So it allows us to connect small sites, but we're building much bigger sites.

Josh Baer: And to clarify, at last earnings, you had line of sight or a target to three gigawatts. Does the announcement from yesterday, is that incremental to that three, or is that a step toward that progress of hitting that target?

Arkady Volozh: This is actually an example of what we -- I think independence was included into our announced -- because the contract there was closed, means everything was signed and binding before approval by the authorities. It's another step, approval. There will be more steps, always it's a long process. And again, and it goes in dozens of different locations. If one of the location is delayed, fine, we have others going, that's how we derisk ourselves.

This stage yesterday was just passing a very important stage of local approvals. So it was included into two gigawatts, but there will be more coming, and we said that we probably will be able to announce at least 3 gigawatts this year. It's a pipeline. It's the table which our guys have hundreds of lines there with sites, potential sites. Of course, the bottom 100 will never materialize, but the top 100 of sites eventually dozens of them will come into life.

Josh Baer: That's helpful. Marc, so we've been talking about the scale of what you're building. How does the demand -- what does the demand look like? And what gives you confidence in that there'll be demand to align with this build-out?

Marc Boroditsky: Well, I'm sure that over the last couple of days, you've heard from a lot of executives that demand is extraordinary right now. Demand far exceeds our supply. The opportunity for us is actually to optimize our model to be able to build a high-margin, durable long-term business to take our abilities to generate expanding pipeline and drive the future capacity decisions that Arkady just mentioned a moment ago is the vision. Today, we have so much demand that we have to actually prioritize. So we're taking the initiative to shift away from a transactional mindset to a value-based and solution selling approach. Our sellers spend time with customers, learning about near-term and long-term requirements, understanding how our capabilities tie to their business opportunities. We use that in order to be able to prioritize customers that actually align against our business model. And as a result, we are now driving sale of future capacity at an increasing rate. We're also delivering longer contract terms. As a matter of fact, the last quarter, we shared that quarter-on-

quarter, the number of year-long contracts doubled. We also saw a 50% expansion in ASP, and we saw prices go up even on older generation GPUs.

Now we're also -- and we haven't shared this yet, but we've added it to our set of priorities, we're now securing upfront payment from customers, up to and including, in some cases, customers paying 100% upfront. So based on this principles-based approach, we're able to prioritize the kinds of customers, the kinds of deals that we're looking for that support the overall business model that we're going for.

Now you asked about demand. What really excites me about demand right now is actually how we grow with our customers. Customers that have landed on our platform to start with a small project and then expand with us over time. Some customers are doubling quarter-on-quarter. They're expanding the use cases that they have on our platform. But what really gets exciting is when they actually move to revenue generation. And then we're actually winning and growing with them at the same time. and that is a durable opportunity for us that actually will extend into the future.

Josh Baer: So when you're talking about securing upfront payment, you're not just talking about Microsoft and Meta, you're talking about the broader customer base?

Marc Boroditsky: Correct.

Josh Baer: Great. And to your point on sort of growing with customers over time, I think that ties pretty closely to training versus inference, what are you seeing as far as the mix on your platform? And how does that evolve going forward?

Marc Boroditsky: Well, first of all, I think we all know this, but I want to say it anyways, for AI to deliver on the commercial vision, inferencing needs to outstrip training period. That's the inferencing at the end of the day, is the commercial realization of the functionality somebody is delivering. Inferencing is growing faster than training on our platform. And we're seeing it along three key vectors: customers that are using our platform to do their own inferencing; customers that are actually inferencing platforms that have been built on Nebius; and then our own inferencing solution, Token Factory.

Interesting vector of demand, I think some of this has been reported in the technical press, is the fact that we're now also seeing software vendors who have built functionality on the foundational models that are looking for improved performance, quality and cost that are now diversifying their models and utilizing open source or even training their own models. And that's opened up opportunity for us to sell our inferencing solution in the software vendor sector. We believe this line of demand is going to continue into the enterprise and be a key way that we're going to potentially be able to service enterprise requirements.

Josh Baer: Excellent. Arkady, you have some really large customer contracts and customers, but you also have this very clear focus on the broader AI stack, for general customers and a broader tail. So how do you balance the two, particularly thinking about the build-out of dedicated large GPU clusters versus servicing this broader market opportunity?

Arkady Volozh: So yes, true. Everybody talks about our big contracts, about Microsoft contract, about Meta contract, billions of dollars, and those are great contracts, and those are great customers. But don't get distracted. Our main business model is not there. We are building a full-scale cloud for the rest of the AI builder market.

Big tech companies are great. There's a lot of experience. The teams respect each other when they work together, but our focus is multi-tenant cloud for the rest. We go into these contracts because there is a huge demand from their side; they need to grow faster than they can, and they outsource this growth to companies like ours and many others. But we used to be on their side before. We understand how they're thinking. Eventually, in several years, they will be building their own capacity. They just cannot cope with the speed now. That's why they're outsourcing.

But we treat this business for us is temporary. We're just helping them to grow now. Why? Because for us, we need to grow our part. To grow our part is hugely capital intense. We need tens of billions of dollars to build it. Where we get it, one of the source in those contracts. The contracts themselves bring us revenues. There's many other terms in the contracts which help us to finance our main part. So today, I would say it's only roughly half and half, half of our capacity is serving these huge customers, but they buy just basic services, bare metals. They have their own stack. They don't need our services, higher than just bare metal.

We serve them. We're happy to have them. I hope we are a good outsourcing for them, but we use this, their resources, cash, credit ratings and other things, which we use, utilize to build our own. And this is our main focus, to build our own capacity, which goes to multi-tenancy to the rest of the market to all these AI start-ups becoming corporations, as they grow, all the enterprise market going into AI and building into their processes. This is our main market. And we're building for them.

So today, yes, we're serving these big clients. There's a great contracts, profitable, allow us to finance very good clients. But don't get distracted, our main business is in the second part. And that part further down will be larger and larger.

Josh Baer: That's very clear. So you have these ARR targets, Marc, for \$7 billion to \$9 billion in ARR exiting 2026. I think the Microsoft and Meta deals get you about halfway. So could we talk a little bit about the other half? Like how much is already booked in your current run rate? And what gives you confidence that, that gap can be -- that you can achieve that through these broader customer relationships?

Marc Boroditsky: So you're right, Josh, that the combination of existing customers and our big deals gets us about halfway. What we're targeting with the execution this year is the other half. And in order to deliver that, my confidence is very high, first and foremost, in the demand that we're seeing, so what we already described earlier. And that's with a very modest market coverage. So what gives me greater confidence about our ability to deliver against that is the incremental results that we're going to generate as we expand our coverage this year. So we are dramatically expanding our overall go-to-market footprint across all the regions.

Third is the expansion that I mentioned earlier, that we experienced with customers as we partner with them. By aligning against their longer-term requirements and selling future capacity, we're actually setting ourselves up based on their journey towards expanded use of our platform.

And then the fourth, which is selling the platform, as an example, selling Token Factory, having gone from helping model creators to now supporting inferencing and then now more and more developer tooling, we've expanded ASP and opened up TAM and increased the routes that we have into accounts. So across those four vectors, we have clear initiatives and programs to be able to deliver the results that I mentioned.

I also want to make sure and be clear, though, this all lies on the fact that we have this big vision of capacity that we're bringing online. And as I mentioned earlier, when we talked about demand, it's really a supply-driven marketplace. So we have a great opportunity with the initiatives that Arkady mentioned a little while ago.

Josh Baer: Excellent. I wanted to come back to a couple of elements on the cost side that you have been speaking to throughout. One, the self-build and the other building your own racks. Could you talk a little bit about the total cost of ownership advantages that you have?

Arkady Volozh: As I said, we are a full stack company, just like hyperscalers. When you talk about the size of the industry, you see the revenues of all the participants, of all the companies of each layer, companies who build revenues to build data centers, revenues to build racks and equipment, revenues for different services on top of it. You see the same revenues actually several times. You need to look

at the margins and like value-added tax. You see where the value should -- is created and how much in each layer?

So in our case, we are working through the layers. It means that on the sites which we are building ourselves versus leasing. Leasing is great because you get it already now and you can grow fast, but you pay their margins. You can design your own racks or you go and buy racks. You can buy racks, fine, it's faster maybe, but you pay the margins on top. The same on each layer. When you build everything through the layers, of course, we save on all these layers. If you have your own capital, another layer actually of value creation because capital is one of the most expensive layers, so to speak, here.

We're happy to start with some cash. We are a public company, we were ready to get money through different instruments. So far, we don't have any asset-backed financing involved. We will, but it's not. Whatever we build now is built with our own resources for now.

So, we save on all these technological layers. We save on capital. And through all this stack, we hope that our TCO, total cost of ownership, is probably not the highest on the market. And plus we built additional services when people buy just compute, it's not bare metal only, sometimes it is, but for other customers. It's not. It's other services, they buy storage, they buy different tools and services on top, and they pay separately. And it all creates layers of additional value, which we added to our revenues.

So we think that building this vertically integrated platforms is what- this is where the competition will be because the major players of this market, the three hyperscalers, they all are vertically integrated. They all outsource, of course, now, but basically, it's all theirs. And we are playing in the same league. And we're trying to play not just vertically in the same league, not only through all the services, but also in scale.

Josh Baer:

Thank you. To round out the conversation talking about some investments, first in go-to-market, and then we'll cover CapEx. So Marc, thinking about go-to-market, where are your priorities? I mean you sort of previewed that in one of your answers. But if you think about your investments looking ahead and building out go-to-market, what- where is your focus area?

Marc Boroditsky:

First and foremost is to expand market coverage. I think I mentioned a moment ago that we've got a modest level of coverage and the opportunity is actually to meet the customers where they are today, with a priority to focus on the AI natives that are out there to win as many of them as possible. But we're going to be building an AI builder community. We already have a substantial one around the company, but our vision is to bring one million builders into our community in the next year. This is a great opportunity for us not only to gain access to AI start-ups, but it's also a way into the entire AI ecosystem. So vendors and solution -- pardon me, integrators and software vendors as well as enterprises. I've seen this play out in prior lives, and it's a phenomenal flywheel.

Speaking about enterprises, the second area of investment is enterprise readiness. As I mentioned earlier that we think the enterprise market is going to kick into gear in the coming years. And we're going to be set up to make sure that we can arm the enterprise sales team that we're hiring now with the tech readiness. So we're building already the table stakes functionality from a security and compliance standpoint as well as arming them with the playbooks that are vertically oriented to support them against workloads that we know that are there today.

The third area of investment is up-leveling our positioning. The team has done a fantastic job of entering the market as disruptors. But based on the success that we've had to date and with our sights set on the enterprise, we're going to reposition against the broader opportunity in the market. And you're going to be seeing things this year that support that.

And then fourth is to be able to make sure that we are set up to sell the entire platform. So making sure that the teams are supporting execution against a broader solution-based sell. So being able to align to our customers' near- and long-term goals and sell the entire software stack.

- Josh Baer: Really helpful. With regard to CapEx, there's a \$16 billion to \$20 billion guide for this year. Could you- you have talked about sort of the sources of funding that CapEx. Could we review? And I'm also wondering if your comments before about securing upfront payments, if that changes the equation at all?
- Marc Boroditsky: It certainly helps. I mean, upfront payment is not only a way to ensure that we have the capital to be able to support the deployment of capacity, it also aligns us with the customers' longer-term requirements. So we're now having a conversation and a commercial arrangement that goes beyond today's transactional requirements.
- In terms of the balance of CapEx requirements, we have a number of different ways that we're pursuing support for the balance of capital that we need this year. I think we shared on the last earnings call that 60% of it is already supported by existing opportunities.
- Josh Baer: Okay. That's helpful. And have you broken out that CapEx for this year as far as like how much is self-build, building future sites versus going straight to GPUs that will be generating revenue near term? Any frameworks for thinking about the mix of CapEx?
- Arkady Volozh: As we said, we have our own sources of capital, which we deployed already. We have this big contracts, which help us to finance the rest of the build. And finally, we still own all of our revenues, which is an important difference since we are not paying back to the banks since we own the stack and all the margins are ours, revenues which we generate is also a source of our financing.
- So I don't know, maybe-
- Marc Boroditsky: We do know what the requirements are. So the- and the allocation between a data center implementation and then the build-out is a lot- just about 20% of the total capital requirements are for the building of the data center and about 80% of the capital requirements are for filling it with GPUs. And we have a series of projects in order to be able to support the capital requirements that we have this year.
- Josh Baer: Excellent. You've recently made an acquisition, I think, Tavily. What are you seeing in agentic AI? And how are you positioning for that? And what should we expect from an M&A perspective going forward?
- Arkady Volozh: Well, Tavily is a specific domain in agentic search. Agentic search is a very demanding function. As you know, up to 20% of all the requests to the models to LLM actually need checking the results on live, web-based, mostly data, and they all utilize some kind of web search, and you need a global index to do it. Until last summer, the major global indexes were available through APIs to many players. It was Google's API, it was Bing's API. But then eventually, they were closed or almost closed. All the big players started developing their own functionality in search. They build their own indexes.
- But the rest of the market, the longer tail, they don't have this access. And having this on our platform to give them -- make them competitive to the big guys in the industry. That's why we want to have this functionality on our platform. It's one of many which we have, but this is very important for today. That's why the Tavily acquisition, but it's not only that. We had some experience with building global indexes before. So we hope we know what we're doing here.
- Josh Baer: Excellent. We're out of time. Arkady and Marc, thank you so much for the conversation.
- Marc Boroditsky: Thank you. Appreciate it.
- Josh Baer: Thanks, Josh.