

# NBIS Earnings Call Transcript

**Date: 2025-11-11**

**Quarter: 3**

Neil Doshi: Thank you, and welcome to Nebius Group's Third Quarter 2025 Earnings Conference Call. I'm Neil Doshi, Vice President of Investor Relations. Joining me today are Arkady Volozh, Founder and CEO, and our broader management team. Our remarks today will include forward-looking statements, which are based on assumptions as of today. Actual results may differ materially as a result of various factors, including those set forth in today's earnings press release and in our report on Form 20-F filed with the SEC. We undertake no obligation to update any forward-looking statements. During this call, we will present both GAAP and certain non-GAAP financial measures. A reconciliation of GAAP to non-GAAP measures is included in today's earnings press release. The earnings press release, shareholder letter, and accompanying investor presentation are available on our website at [nebius.com/investorhub](https://nebius.com/investorhub). And now I'd like to turn the call over to Arkady.

Arkady Volozh: Thanks, Neil, and thank you, everyone, for joining the call today. I'd like to share my thoughts about the demand environment, our capacity plans, and what we're doing in our product. First, about the demand. Q3 demand was very strong. We sold out all of our available capacity. We continue to see a consistent trend: we bring capacity online, we sell all of it. With the new generation of NVIDIA Blackwells coming online, more customers are interested in purchasing capacity in advance and securing it for a longer period of time. Today, we're very pleased to announce that we signed another major deal, this time with Meta, for approximately \$3 billion over the next five years. In fact, demand for this capacity was overwhelming, and the size of the contract was limited to the amount of capacity that we had available. Which means that if we had more, we could have sold more. This deal comes on top of the Microsoft deal we announced early September, with a contract value between \$17.4 billion and \$19.4 billion. As we said before, we expect to sign more of these large long-term deals, and we're delivering on that promise. As busy as we are with these mega deals, our main focus is still to build our own core AI cloud business. We made great progress here with AI-native startups like Coursera, Black Forest Labs, and others. The economics and the cash flow of mega deals are attractive in their own right, but they also enable us to build our core AI cloud business faster. This is our real future opportunity. Now on the capacity. In order to meet the growing demand, we have accelerated our plans to secure more capacity, and this is actually our main focus for now. Capacity today is the main bottleneck to revenue growth, and we are now working to remove this bottleneck. As we look to 2026, we expect our contracted capacity to grow to 2.5 gigawatts. This is up from the one gigawatt which we discussed in our previous earnings call in August. Furthermore, we plan to have power connected to our data centers, which means fully built, of approximately 800 megawatts to one gigawatt by the end of 2026. While we made significant investments in our capacity footprint, we're also investing in our main product, our AI cloud. To expand our addressable market opportunity to large enterprise customers, we released our new enterprise-ready cloud platform version 3.0 called Ether, and our new inference platform called Nebius Talking Factory. We believe Ether gives organizations the trust, control, and simplicity they need to run their most critical AI workloads. Nebius Talking Factory is an inference platform that enables organizations to run open-source models with reliability, visibility, and control. And we have a large pipeline of new software and services that we are continuing to build, which will differentiate us from other cloud companies. Based on the strength of demand that we see and our accelerated capacity growth plan, we believe we can achieve annualized run rate revenue

(ARR) of \$7 to \$9 billion by the end of 2026. In summary, Nebius is positioned to win in this large and rapidly expanding AI cloud market. We're just beginning to realize the powerful potential of the AI evolution that is underway, and we are quickly becoming one of the primary cloud and infrastructure providers to support it. And with this, I would like to hand the call over to our CFO, Dado Alonso. Dado, please.

Dado Alonso: Thank you, Arkady. While the details of our Q3 financial performance can be found in our shareholder letter, I'd like to provide some additional color to the quarter, discuss our financing options, and conclude with 2025 guidance. Q3 group revenue was \$146 million, up nearly 355% year over year and 39% quarter over quarter. Annualized run rate revenue for the core business at the end of September was \$551 million. The core infrastructure business, which accounted for nearly 90% of total revenue, grew 400% year over year and 40% sequentially. Once again, we sold out our capacity, and our revenue growth was limited only by the capacity that we were able to bring online. I'm also pleased to say that adjusted EBITDA margin for the core infrastructure business expanded quarter over quarter to nearly 19%. On financing, in order to support our aggressive growth plans in 2026 and to maintain this pace of growth in 2027, we will be utilizing at least three sources: corporate debt, asset-backed financing, and equity. We are in the process of raising asset-backed debt, which we'll be able to secure with attractive terms supported by the creditworthiness of our largest customers. Tomorrow, November 12, we will be putting in place an at-the-market equity program for up to 25 million Class A shares and plan to file a prospectus supplement. We will evaluate the program regularly based on our capital needs. The program enables us to access equity funding on an efficient ongoing basis. However, we will remain dilution sensitive as we prepare to finance future growth opportunities. Now I would like to turn to 2025 guidance. As we approach the end of the year, we are tightening our full-year group revenue guidance to a range of \$500 million to \$550 million, and we are currently pacing to the midpoint of that range. This compares to the \$450 million to \$630 million in our previous guidance. The reason we are in the middle and not at the top of that range simply relates to the exact timing of when capacity comes online. Our current momentum and long-term trajectory remain extremely strong. Our annual run rate revenue, which is a good reflection of our future growth opportunity, continues to expand, demonstrating the resilience and scalability of our business model. As such, we remain well on track to hit our ARR guidance of \$900 million to \$1.1 billion by the end of 2025 while also paving the way for substantial annualized run rate revenue growth in 2026 and beyond. In terms of the mega deals, we will begin serving Microsoft and Meta late in the quarter, and almost all of the revenue from these deals will start to be realized and ramp up during the course of 2026. We plan to give full-year revenue guidance for 2026 next quarter. Turning to adjusted EBITDA, as we have previously indicated, we expect to be slightly positive at the group level by year-end while remaining negative for the full year. Regarding CapEx, we are raising our 2025 guidance from approximately \$2 billion to circa \$5 billion. This acceleration reflects our strong conviction in the demand outlook and our decision to secure critical infrastructure, including hardware, power, land, and key sites. These investments are strategic enablers of future growth and will position us exceptionally well to capture the opportunities ahead. In summary, we have a large and rapidly growing opportunity in front of us, and we are executing with focus and discipline to capture it while delivering substantial sustainable growth and setting the stage for strong long-term profitability. Now let me turn the call over to Neil for Q&A.;

Neil Doshi: Great. Thank you, Dado. We'll give it a moment to collect questions from the online platform, and then we'll begin the Q&A.; Alright. Let's start with the first question coming from Alex Platt of D.A. Davidson. Can you tell us more about the new Meta deal? Why did they choose you, and how should we model the deal?

Arkady Volozh: Well, again, as we're happy to announce today, this new deal with Meta is approximately \$3 billion. As I said, the size of the deal was limited only by the capacity that we had available, and if we had more capacity, we could have signed a bigger deal, probably. After we announced Microsoft in September, we said that we would have more deals of this kind, more large deals. And actually, we're delivering on that promise, and we're optimistic as these deals will arise more and more. However, these mega deals are important, but it's important to stress that we will remain focused on developing our own AI cloud, which currently serves not only these big deals but AI startups and enterprises. Ultimately, we believe that these large contracts provide us with great sourcing of

financing for us to continue building our core AI cloud business.

Neil Doshi: Great. Thank you, Arkady. We'll take the next question from Alex Duval, our analyst from Goldman Sachs. So we provided the updated 2026 ARR outlook of \$7 to \$9 billion. What exactly is in the \$7 to \$9 billion ARR target? And is this based on preexisting core business plus Microsoft and Meta? Is there anything else in terms of signing up for large deals? Mark, maybe you can take this one.

Marc D. Boroditsky: Thank you, Alex, for your question. Let me walk you through the building blocks of how we get to this \$7 to \$9 billion in ARR. First of all, as we already shared, we had a bottleneck in capacity, and we worked extremely hard over the last several months to unblock this bottleneck. As we shared, we plan to have 800 megawatts to one gigawatt of connected power by 2026 and 2.5 gigawatts of contracted power. Second, we see the demand out there from AI startups to enterprises to the large strategics. And we see that client demand that we were unable to sell this past year due to a lack of capacity, and we strongly believe that the capacity we are putting in place in 2026 will help us to meet more of this demand. At the end of the day, we will allocate between the categories of customers based on the individual economics of the deals they represent. Thirdly, this new capacity that we are putting in place together with our current capacity that has already been sold, and the long-term contracts that we signed with Microsoft and Meta give us the confidence that we can achieve the \$7 to \$9 billion of ARR, of which more than half is already booked.

Neil Doshi: Right. Thank you, Mark. We'll take another question from Alex at Goldman. Can you walk us through the timeline of your infrastructure build-outs for Q4 2025 and 2026? And what gives you confidence that you can reach your 2.5 gigawatts goal for contracted capacity? Andrey?

Andrey Korolenko: Hello. Yeah. Thanks, Neil. So we are ramping up our capacity as fast as we can to accelerate our growth for the next year and beyond. We are happy to launch now already in Israel and the UK, and all the capacity in those regions was presold before the launch. And we are growing with the number of regions where we are present, and we are also bringing new capacity online in the current sites. We're also coming online in New Jersey. We are launching new phases of Finland in Q4, which are also presold, by the way. In 2026, we will continue scaling the existing data centers, including the UK, Israel, and New Jersey, and we have new data centers already in development, both in the US and Europe, and they start to come online in 2026. We are also in the process of securing several new large sites, which we believe will add hundreds of megawatts, and some of those will go online by 2026. So overall, at the moment, we are looking at more than or around 2.5 gigawatts of contracted power by 2026. And as we said, demand is growing massively, and we are very focused on rapidly building the capacity and the future pipeline to meet the demand in 2026 and beyond.

Neil Doshi: Great. Thank you, Andrey. Alright. We'll take a question from the folks who've been submitting questions. We're getting a lot of questions on Microsoft and Meta revenue. How should we be thinking about revenue contribution from Microsoft and Meta deals for this year and going forward? Dado?

Dado Alonso: Well, the Microsoft contract will not have a material effect on our revenue and ARR in 2025, as the first tranche was just delivered. All of our remaining tranches will be delivered in 2026, with more than half of them during the second half. So we actually expect revenue to ramp up over the course of the year. Starting in 2027, we will begin to recognize the full annual revenue run rate of this new Microsoft deal. With regards to Meta, we will be concluding the deployments within the next three months, so we expect to mostly be at a full revenue run rate in 2026.

Neil Doshi: Great. Thank you, Dado. Maybe another question from our online audience. What does the overall demand environment look like in Q4 and into the next year? Mark, do you want to take this?

Marc D. Boroditsky: Certainly. Certainly. I joined the company about five months ago, and I've had an extraordinary experience in these past five months. It's extraordinary from the standpoint that I've never seen the kind of demand profile that we're experiencing. It is literally accelerating for Nebius, and I believe as well for the broader market. As an example, in the recent quarter and quarters before, we saw pipeline generation. This is opportunities by customers that want to buy from us. Expand. As a matter of fact, in the past quarter, Q3, we saw pipeline generation expand 70% quarter on quarter, and we generated \$4 billion in pipeline in that quarter. But we were only able to convert a portion of that given the constraints of our capacity. As a matter of fact, I've learned a new skill, one I don't think many go-to-market professionals have ever had to experience, and that's learning to say no to customers. As

we routinely sell out and have to actually let them down lightly and try to convince them to purchase in the future. As I look out to 2026, and I think through the demand profile, the kind of pipeline that we're generating right now is giving us high confidence to continue to expand our results and drive towards the ARR growth that Arkady mentioned earlier on the call.

Neil Doshi: Great. Thanks, Mark. We have a question now from one of our analysts, Nehal Chokshi from Northland. Incremental ARR in September was around \$12 million, down from \$180 million the prior quarter and \$159 million in March. Why is incremental ARR down? Neil, it's a great question. You know, as we've stated, a lot of our revenue and our ARR is really dependent on us being able to bring on capacity. And because capacity really has been the bottleneck, that's why we've seen a little bit of that trend. However, as we're bringing on a lot of capacity in Q4, you should see that incremental ARR in Q4 will be significantly higher. Alright. Let's go to another question from online in terms of the CapEx. You have just announced your plan to achieve connected capacity of 800 megawatts to one gigawatt by 2026. How are you thinking about CapEx, and what is your philosophy on CapEx spending?

Arkady Volozh: I think Arkady, can I take this? Probably, I should take it. Again and again, as we see, at least this year, our revenue growth was limited by our capacity, and everything we built was ultimately sold. So in theory, we should try to build as much as we can. In practice, we are limited by certain physical world limitations. The physical world cannot grow five or ten times a year. We have limitations in the supply chain and obtaining permits. The amount of capital that we can deploy. So when we plan for data center CapEx, there are actually three stages there. The first stage is securing the land and power. The second stage is building the data centers themselves, doing shell, electrical and cooling equipment, batteries, and so on. Physical installation, which we call connected power. And the third part is finally deploying the GPUs themselves. And if we look at it from the CapEx point of view, speaking, it breaks into three spending blocks. So first, stage, securing land and power. It's pretty cheap. It's around, again, it depends on the scale, but it's around 1% of total CapEx for securing those ports and electricity. The second stage, building with this set of building connected power, is something around, I don't know, 18, 20%. And the remaining 80%, the main part, is for deploying the actual GPUs. This is the main part of CapEx. So if we want to build as much as our capital will allow us, what should we do? First, we should secure as much capacity as we can because the cost, actually, it's not so it's an immaterial at this scale. Second, we should build as much as our capital allows. And third, we will fill GPUs in line with contracted or clearly visible demand. We will need this massive 80% spend will come only when we see real demand. That's why we say that in 2026, we will be securing 2.5 gigawatts total contract capacity, and we are planning to physically build 800 to one gigawatt of connected data centers. This will be done by the end of next year.

Neil Doshi: Great. Thank you, Arkady. Let's see. Another question from Alex Duval from Goldman Sachs. You have announced your target is 2.5 gigawatts of contracted power, whereas before it was one gigawatt. Is it fair to assume that if you get 2.5 gigawatts, this will equate to over \$20 billion of revenue? By when do you envisage you could do this and how? Maybe we'll give this to Andrey.

Andrey Korolenko: Thank you, Neil. I guess it's fair to assume, but as Arkady just mentioned, we are securing the access to the power and the ability to build. But we will invest CapEx actually in building out and deploying the GPU in those. Keep in mind the constraints that we have with the capital and according to the demand in the future periods. Just important that we are able to accelerate when it will be needed, so we don't like being blocked by the capacity constraints all the time.

Neil Doshi: Great. Alright. Take another question from online. Is it in a situation when you are sold out, is that the same issue? Or is that really an issue with your future growth and differentiation of servicing a broader range of customers? Mark, can you take this one?

Marc D. Boroditsky: Certainly. Thank you, Neil. It's a great question. I mean, in theory, the situation of being sold out is a nice problem to have. But the person asking the question is right. For our business model, it's really important for us to be able to not only service large tech companies but also be able to support our AI cloud and in a very diverse set of customers. As a matter of fact, servicing startups and software vendors and enterprises is not only about delivering on their capacity needs today. We want to build partnerships with these customers and help them to meet their capacity requirements in the future, especially with enterprises. Because they don't want to actually have a multitude of vendors. They prefer to align with a strategic partner. That's why we are working very closely with Andrey. And

as Andrey mentioned earlier, as we look forward and think about deploying capacity, it's going to be based on the demand that we're seeing out there. So utilizing the pipeline that we're building and the demand that we're experiencing to work with Andrey to identify the capacity that you should deploy. It's a very dynamic model that we're trying to put in place.

Neil Doshi: Great. Thank you, Mark. Appreciate that. We have a question from Nehal Chokshi from Northland. Who's asking, you know, going, you know, so you've done equity deals. We've also done, you know, equity-linked deals as well. Dado, you know, how are we focusing on debt and asset-backed financing for large deals?

Dado Alonso: Thank you, Nehal. Well, as you know, this is a capital-intensive business. And as we've said previously, funding our growth will require raising a significant amount of capital. In this context, we are actively evaluating a range of financing options today, including asset-backed financing, corporate-level debt, and equity financing. And we are working on all fronts in order to maintain a disciplined capital structure and maximize our shareholder value. With regards to asset-backed financing, we believe that we will be able to secure such a facility with attractive terms supported by the creditworthiness of our largest customers. I would like to reiterate that as we are growing our business, our focus and ultimate goal is to maximize our shareholder value.

Neil Doshi: Great. Thanks, Dado. And maybe just sticking on the theme of financing, from the online portal. You know, why are you planning to pursue an ATM? You just completed a secondary, and this will result in additional dilution to shareholders.

Dado Alonso: Any share of some perspective. We will be putting in place an at-the-market equity program for up to 25 million of Class A shares, and we plan to file the prospectus supplement tomorrow. We want to make sure that we have more tools at our disposal to access capital markets when needed. This is a long-lasting program, which will be used along with other capital raise options, including corporate debt, asset-backed financing, and others, as I mentioned in my opening remarks and just before this question. So the program enables us to equity funding on an efficient ongoing basis. However, we will remain dilution sensitive as we seek to finance future growth opportunities.

Neil Doshi: Great. Thanks, Dado. Let's see. Another question from online. How are the early operations of your new UK facility progressing? Tom?

Tom Blackwell: Yeah. No. Absolutely. So short answer is progressing very well. You may have seen just actually last week, Arkady and a few of us were there. We had our official launch as we presented the data center to the UK market. This is actually capacity that will be coming online really actually in the next week or so. So pretty very in the coming days. You might remember actually that in June was when we first announced our intention to launch in the UK, and actually even in the time since June and now, we've already come close to doubling the capacity that we're bringing on stream. You know? And that's just really a function of extremely strong demand that we're seeing in the UK. And actually, as is often the case with new that we bring on even before going live, we're pretty much sold out. So if not already fully sold out with our capacity. So that track, you know, that's a trend that just continues. I would just say overall, a few words about the UK, actually. I mean, we're very bullish actually about the opportunities in the UK. It's a vibrant AI market. You know, it's probably one of the most dynamic that we see outside of the US and China. The government's making a big push to support the growth of the industry and having a reasonable degree of success in this field. So there's a we see a lot of AI startups. The VC environment is strong. You also see some of the large tech companies establishing regional R&D; and presence there. So really, there's a lot happening in the UK, and we think a lot still to come for Nebius in the UK. We're very happy to be there. And, actually, although this specific facility that we have, I think, with the capacity once by January will have reached the peak capacity there, we see a lot of other opportunities to expand capacity in the UK overall.

Neil Doshi: Great. Let's see. In terms of we'll take another question on capacity. So, you know, you mentioned this quarter that you're fully sold out of available capacity. What are your constraints to growing in the near term and medium term to capture more of that demand? And could you also address some of the recent comments on the market around power equipment constraints?

Andrey Korolenko: Yeah. I'll take it. Yeah. Thanks, Neil. As we discussed in most of the previous questions, capacity remains our main bottleneck. Everything we deploy, we sell. And we see the demand that continues to significantly outstrip our supply each time we add new clusters. So in the

near term, the key challenges in increasing capacity are securing power and the supply chain. And we're addressing this. We have managed these situations in the past and have quite a bit of expertise in building on the data centers and the duration of those. So overall, we have mentioned. Generally speaking, we are doing quite well, actually, with the pipeline. And when we spoke last quarter, I believe that we announced that we have secured the roadmap of the one gigawatt of power. Now we are talking about the number 2.5 gigawatts. And we're still putting a lot of focus on growing this number and making this number reliable and effective and actually bigger.

Neil Doshi: Great. Thank you, Andrey. Online, we're getting just a few questions about any updates on the New Jersey facility. Andrey, do you want to take that?

Andrey Korolenko: Yeah. The New Jersey facility goes as planned. And the first tranche already was handed over to Microsoft, and we have continued the further expansion to the crane.

Neil Doshi: Alright. Looks like a question from online. Maybe more of a market question. Are you concerned that we are in an AI bubble, Arkady?

Arkady Volozh: Don't have anybody ask this question these days. Well, what we see today, the demand is here. Right? Yeah. We understand that we are in the center of a one-second generation AI evolution. Much no doubt that much more compute will be needed and much more will be built. This situation of unbalanced demand supply is temporary. Of course, eventually, demand supply will level up. And what we are doing in addition just to growing this raw capacity, we are building our AI cloud, which will real businesses, real industries, real enterprise market. Where AI will be making creating value. And we are big believers that the AI industry in general, you know, sector specifically. It's gonna be okay. Ultimately, we just need to be sure that a, we're diversified in terms of customers and workloads. And this is actually what our software is basically doing. That we invest conservatively and that we finance our growth responsibly, and we are very much focused on this. And also, while we're growing rapidly, five times more a year, we still remain laser-focused on maintaining healthy margins. And a sustainable business model as a whole. In old days, I would say. Healthy unit economics. So we are focused on that. And whatever comes will come. I would hope that we will be okay.

Neil Doshi: Great. Thank you, Arkady. Next question is from Alex Platt from D.A. Davidson. How should we think about the lead time between when power is connected to and when it is hooked up to GPUs and generating revenue?

Andrey Korolenko: Yeah. Thanks, Neil. So on the technical side, it also depends if it's a new site or if it's a special one of a current site. But, generally speaking, from the connected power and start of the GPU deployment until it can go in the platform engineering revenues anywhere from six to twelve weeks. If it's already an existing site, that can be even quicker. But generally, we also have flexibility. That's why we are building infrastructure. We have flexibility when we deploy and when we deploy how much we deploy.

Neil Doshi: Great. Thanks, Andrey. Alright. Question from online. Can you update us on your progress with your primary customer segments? Mark, can you help us with this?

Marc D. Boroditsky: Certainly. Certainly. We continue to see extremely strong demand from our customers in our core AI business, and we're continuing to expand business overall with our existing customers. As a matter of fact, we added a number of new customers in Q3. Most notably, some very disruptive startups like Cursor AI, Black Forest Labs, and World Labs. I'm sure everybody's heard of Cursor. We're very proud to be their partner. For those that haven't, they're an extraordinarily popular AI-powered code editor that is helping millions of developers to write and debug and optimize their code faster, and they're making great strides into the enterprise. Black Forest Labs is an interesting customer that is developing cutting-edge generative AI models specifically for image and video generation. Their popular Flex One model helps turn text and images into high-quality media-ready visuals. And World Labs is building something they call a large world model, which is able to simulate 3D worlds and it gives developers and AI engineers the necessary spatial awareness to build applications for things like media and gaming and architectural design and as well for physical AI and robotics. We've also, as I mentioned, seen expansion with existing customers. As a matter of fact, as an example, we've seen expansion with our software vendor customers like Shopify, and then also we've made great strides with our efforts around driving vertical market success, adding significant customers in our healthcare life sciences part of the business. And we're also making significant advancements in physical AI and

media and entertainment customer segments.

Neil Doshi: Right. Thank you, Mark. Maybe a question to look like this is a question for Dado. Few people are asking, any puts and takes that you can provide on your revised 2025 year-end revenue guidance?

Dado Alonso: Well, our business is actually scaling rapidly. There can always be fluctuations in the exact timing of deployments in such a fast-growing company like ours. And our focus remains on building a very large company. Obviously, much larger than today and significantly bigger than our plans for 2026. This was and continues to be our main focus. In any event, our annualized run rate revenue, which is a better reflection of our future growth opportunity, continues to expand, demonstrating the resilience and scalability of our business model. As such, we remain well on track to hit our ARR guidance range of \$900 million to \$1.1 billion by the end of 2025 while also paving the way for substantial revenue growth in 2026 and beyond.

Neil Doshi: Great. Thank you, Dado. See, a question from the online community. How is your enterprise initiative ramping up? We seem to make some good improvements there over the past couple of quarters. Mark, do you want to help take this?

Marc D. Boroditsky: Certainly. Certainly. We are making strides with regard to becoming enterprise-ready. As you saw with the launch of Nebius 3.0, what we call Aether, we've delivered a number of AI cloud improvements to support enterprise requirements. As an example, in the release, we are delivering really important compliance and security certifications. And we did this, as a matter of fact, in a matter of months when it would normally take other organizations a lot longer to deliver these types of capabilities. As a matter of fact, we delivered some important functionality that enables enterprise administrators to proactively manage their implementation. So tooling and controls like identity and access management and dashboards for evaluating the performance and security of their AI. I think as we all know, the sort of the critical foundation for enterprise readiness is to have these kinds of compliance and security certifications in place. And the enterprise functionality that enterprises are looking for. And the third is to have an enterprise-ready sales team. On that front, we are adding a number of key leaders to our organization. And we are expanding the overall sales organization for coverage in enterprise software vendors and key verticals. It'll take some time for the sales team to ramp, but we are building the foundation between the functionality that I mentioned and the overall team coverage. That I think will set us up for a strong 2026 with enterprises.

Neil Doshi: Great. Thank you, Mark. Keeping with our online investor base, you know, you recently launched Token Factory. What is the opportunity around this? And does this expand your market or open up new segments? Maybe we'll ask Roman to take this one.

Roman Chernin: Yeah. Thank you, Neil. Happy to talk about our new launch. So I will start a little bit from demand evolution. We fairly see now the next wave of AI demand growth, and it's mostly driven by the companies, by the people who apply AI to real-world applications across all industries, in B2C and B2B. It's not necessarily foundational model builders like it was, let's call it, in the first wave. And we as Nebius realized that we needed our inference as a service offering to make it serve a broader set of customers, including enterprises. So Token Factory gives vertical AI product builders, ISVs, and enterprises a platform to build what we call a flywheel of applying LLMs in vertical AI use cases at scale. Transforming, we help them to transform open-source models into optimized production-ready systems with guaranteed performance and transparent cost per token. We obviously leverage the underlying to bring the most efficient, scalable solution to our customers when they can be sure that they get the best total cost of ownership and can confidently grow with us. As a result, organizations can deploy and scale models such as OpenAI OSS, Qwen, Deepsea, Glamour, Nematron, and many others on dedicated endpoints with guaranteed performance tuned for the super latency and 99.9% uptime. So in total, I mean, I must say we are excited about the opportunity of inference workloads. We believe that all companies will invest in inference to productize AI, and for us, it means it will require significantly more compute and will support this wave of growth. As well as we do for foundation model builders.

Neil Doshi: Great. Thank you, Roman. Jumping back to online, looks like we have some additional questions here. What demand are you seeing for the new Blackwell generation, and how is this demand from the previous Hopper generation? Mark, do you want to take this one?

Marc D. Boroditsky: Yes. Certainly. Thank you, Neil. Demand remains very strong across all types of GPUs. And as we said, we sold out our capacity in Q3, and that's across all types. And we're nearly sold out with respect to Q4. Talking about the Hoppers, we continue to see extremely positive demand for these chips. An interesting set of dynamics that we're experiencing is that as customers come to their renewal for Hoppers, or if they're looking to upgrade to say Blackwells, in both cases, we're typically selling them immediately. And often at better pricing than they were previously priced. As we're actually in tandem rolling out the Blackwell. So very strong demand profile for existing Hoppers. We're also seeing very strong demand for Blackwells. We're benefiting from the fact that we're one of the first companies to deliver them in the market. In our Israel data center, we launched with B200s, and in our UK data center, we launched with B300s. And we've essentially presold much of that capacity before these facilities even opened. We're very excited as well that we're launching GB300s. We're the first to do so in Europe, which will be coming online in our Finnish data center later this quarter in December. Thinking about in-production capacity, but we're also now pre-selling. We're right now, as I mentioned, selling the remnants of Q4. New capacity is being delivered in future quarters. So we're seeing a very strong demand across all types of GPUs. And as I mentioned earlier, and as Andrey mentioned earlier, we're working in close partnership with Andrey's team to make sure that our sales pipeline allows us to drive our model in order to be able to support GPU requirements in subsequent quarters.

Neil Doshi: Perfect. Thank you, Mark. Another question from Alex Platt from D.A. Davidson. He's asking about our strategy regarding larger deals. Do we have a medium-term capacity target for these deals and customers?

Arkady Volozh: Yes. We are very opportunistic here. The demand is there, not only for everyday deals but for large mega deals as well. And we'll enter into the deals which provide us with the best margins. We're very much focused on margins and profitability, not only on growth itself, and all our decisions actually direct from there.

Neil Doshi: Great. Thank you, Arkady. Next question from Andrew Beal from Aarte. Can you provide more details regarding some of the greenfield sites? Do you have LOIs for further new US and EU DC locations, or are you further down the road with these?

Andrey Korolenko: Yeah. Thanks, Neil. Generally, we are making great progress to bring our capacity. We have a robust pipeline both in Europe and the US. We mentioned that we are on the way of securing 2.5 gigawatts of, well, roadmap for the power through the next year. We are in the station also further down the road as well. But we are not in a position to say more at this stage.

Neil Doshi: Great. Alright. Question from online. Can you provide an update on your facility in Israel? Tom?

Tom Blackwell: Yeah. Sure. So as you can see, we're growing rapidly. So just last week, we were in the UK launching, and just a couple of weeks before that, we were in Israel. And, actually, the data center facility that we have there is already fully live. And as I think we've made various references to, Mark's mentioned it previously, again, that was capacity that was effectively even presold, and we definitely have opportunities to expand further in terms of capacity. We think Israel is a great market. You know, we again, we see a lot of demand. There's a lot happening in tech and in AI. And, actually, one of the things that's interesting about the market, I mean, our decision to go in there was purely based on our own commercial considerations. We think there's a great there's a lot of growth for that. But, actually, the government's also doing some interesting things. Really, to stimulate further demand. And so they're actually effectively putting money to subsidize sort of AI startups and institutions and helping them to access the compute as a way of getting of having the growth move faster. So we think it's anyway, we're doing great there. We have a we think that there's a great opportunity for us. And, actually, the model there might even be a model that we think either in other countries might look at that are thinking about building up the domestic demand in the AI industry. So overall, going great.

Neil Doshi: Great. Thanks, Tom. Alright. From our online platform, let's see. How do you think about partnering with or buying potential companies that already have secured power or land? Consolidating or consolidating other neo clouds. Arkady?

Arkady Volozh: Companies, we secure power and land. Again and again, it's all about margins. We are pretty much focused on the margins when we enter into a new contract, when we're raising capital,

when we're developing new plots for data centers, when we design those data centers, when we build our RAC software, and so we vertically integrate and we are looking for efficiency at each stage. We have looked into potential acquisitions of power land market, but so far, our approach process proved to bring, I would say, much higher margins. So far. Yeah. So we're still moving further and further into building our own facilities. And we're actually decreasing the share of our own over-corrugated rented facilities. More and more will be our own facilities. Obviously, we will continue to consider different opportunities, but as you can see, we were able to secure a significant level of contract power organically. So we strongly believe that rather than later, the margins for the infrastructure business will play a significant role in the ability to, I mean, the business to grow and develop. And we basically remain very focused on that.

Neil Doshi: Great. Thank you, Arkady. Alright. Let's take another question. Kind of market-related question. Maybe this one will go to you, Arkady. Is there any chance that GPUs are oversupplied in the coming year as new suppliers come to the market?

Arkady Volozh: Two things here. First, we strongly believe that the market will still be supply constrained, at least in 2026. This means that data center capacity will be the chokepoint. Also, as we mentioned earlier, we plan our capital spend in those three stages: land power, building the facilities, and then the GPUs. And this conservative stage approach keeps us from overspending, actually, and allows us to maintain a healthy financial position. If there are any changes to the market, we'll be in good shape to weather a downturn, we hope.

Neil Doshi: Great. Thank you, Arkady. Let's see. A couple of our analysts are asking, you know, in terms of, you know, any big challenges regarding the completion of the Vineland facility? You know, any, you know, challenges to meeting, you know, any performance obligations of your deal. Thanks, Neil. I think I already spoke about it.

Andrey Korolenko: So yeah, as of today, it goes as planned, and we already handed over the first tranche to Microsoft. So yep, continue working according to the plan.

Neil Doshi: Great. Thank you. Alright. I think we'll end the call there. Thank you, everyone, for joining, and we will speak to you again next quarter. Thank you.