

DigitalOcean Holdings Inc

Q3 2025 Earnings Call

Wed Nov 05 2025 Earnings Call Transcript

Q3 2025 Earnings Call *(Corrected version)*

Summary

NEW

Key Takeaways

DigitalOcean Holdings Inc's Q3 2025 earnings call highlighted strong financial performance, driven by growth in AI-native customers and digital-native enterprise clients. The company exceeded Q3 guidance, raised its 2025 and 2026 outlook, and is increasing investments in data centers and GPU capacity. Key themes included revenue growth, profitability, AI traction, customer expansion, and strategic investments.

Key Financial Results

- Q3 revenue reached \$230 million , a 16% year-over-year increase , marking the highest growth since Q3 2023.
- The company delivered its highest organic incremental ARR in its history at \$44 million .
- Adjusted EBITDA was \$100 million , a 15% increase year-over-year , with an adjusted EBITDA margin of 43% .
- Non-GAAP diluted net income per share was \$0.54 , a 4% increase year-over-year .
- Adjusted free cash flow was \$85 million , or 37% of revenue .
- Trailing 12-month adjusted free cash flow margin was 21% .
- Unlevered adjusted free cash flow was 37% of revenue .

Business Segment Results

- Direct AI revenue more than doubled year-over-year for the fifth consecutive quarter .
- General-purpose cloud products saw the highest incremental organic ARR since Q2 2022 .
- Revenue from customers with more than \$100,000 in annual run rate grew 41% year-over-year , representing 26% of total revenue .
- Customers with greater than \$500,000 and greater than \$1 million in annualized run rate grew revenue 55% and 72% , respectively.

Capital Allocation

- The company repurchased the majority of its 2026 convert in the quarter, strengthening its balance sheet.
- Repurchased \$2.9 million of shares in Q3 , buying back approximately 101,000 shares .
- Cumulative share repurchases since IPO totaled \$1.6 billion and 34.9 million shares through September 30, 2025.
- Entered into an equipment financing arrangement with a third-party financial institution for \$28 million to better align investments with future revenue.
- Secured around 30 megawatts of incremental data center capacity to support growth in 2026 and beyond.

Industry Trends and Dynamics

- Increasing momentum with AI-native customers in inference mode.
- Growing traction with high-spending digital-native enterprise customers .
- Continued strength in revenue from new customers .
- Demand for the company's agentic cloud has exceeded supply.
- The future of AI is in agents and agentic workflows .

Competitive Landscape

- DigitalOcean's unified agentic cloud platform is one of the most intuitive and complete platforms for taking AI agents from prototype to production.
- The company's platform is gaining traction over hyperscalers .
- DigitalOcean's platform offers simplicity, scalability, and reach without the hyperscale complexity .

Growth Opportunities and Strategies

- Continued innovation in the company's comprehensive agentic cloud to support scaling AI and digital-native enterprise customers .
- Augmenting the product-led growth engine with a focused direct sales motion .
- Expanding data center footprint to meet growing demand.
- Strategic partnership with fal.ai to accelerate generative AI content creation.
- Launch of the DigitalOcean AI Partner Program to empower AI and digital-native enterprises.
- Introduction of new features like Spaces Cold Storage and automated storage auto scaling for managed databases.
- Ramping engineering resources to accelerate the unified agentic cloud roadmap.
- Continued targeted investment in new sales and marketing initiatives.

Financial Guidance and Outlook

- Raising full-year 2025 guidance on both revenue and margin.
- Expects to achieve the 18% to 20% 2027 revenue growth target in 2026 , a full year earlier than projected.
- Q4 2025 revenue is expected to be in the range of \$237 million to \$238 million , approximately 16% year-over-year growth .
- Full-year 2025 revenue is projected at \$896 million to \$897 million , representing approximately 15% year-over-year growth .

- Adjusted EBITDA margins for Q4 2025 are expected to be in the range of 38.5% to 39.5% , with approximately 41% for the full year.
- Non-GAAP diluted earnings per share for Q4 2025 are expected to be \$0.35 to \$0.40 .
- Full-year 2025 non-GAAP diluted earnings per share are expected to be \$2.00 to \$2.05 .
- Projected full-year adjusted free cash flow margin of 18% to 19% .
- Expects to comfortably deliver 18% to 20% growth in 2026 .
- Anticipates delivering high-30s to 40% adjusted EBITDA margins in 2026, while maintaining mid-to-high teens adjusted free cash flow .

Q & A Highlights

The earnings call Q&A session focused on several key themes: the nature and drivers of the **eight-figure committed contracts**, the impact of **hyperscaler outages** on migrations, **capacity planning** in light of increasing demand, the company's **competitive strategy** and product development, the potential inclusion of **AI revenue** in net dollar retention (NDR) metrics, and the factors influencing **NDR** and expansion activity, and the outlook for **capital expenditure** and equipment leasing. Analysts sought clarity on the composition of new contracts, the sustainability of AI revenue growth, and the company's approach to balancing growth with profitability.

Eight-Figure Committed Contracts

- **Q: Can you provide more details about the cohort of customers signing these eight-figure committed contracts, including their overlap with AI revenue and any involvement with the private equity community?**
- A: The **eight-figure commitment contracts** come in different forms, primarily from **AI-native companies** using our infrastructure and another customer using our **AI platform** to build agentic experiences for software engineering. It's becoming increasingly difficult to separate **AI** from core cloud, as customers starting with **DigitalOcean** for **AI** are using storage, **VPC capabilities**, and networking. The crossover between the **AI** side and cloud is growing, leading to a more unified cloud platform called the **agentic cloud**. These contracts typically start with **AI** but extend to our **AI cloud** side. Some customers begin with a small proof-of-concept and are now scaling, which is why we're expanding our data center footprint. The **AI workloads** are predominantly inferencing, making them durable and predictable, with opportunities to scale globally. We're focusing on capacity to support marquee **AI-native companies** finding traction with real end customers.

Impact of Hyperscaler Outages and Migrations

- **Q: Are the recent outages at AWS and Azure having any near-term impact or catalyzing more migrations from the hyperscalers, driving more traction for Partner Network Connect or other multi-cloud offerings, and how many of those eight-figure deals were migrations from the hyperscalers?**
- A: We've seen a steady increase in migration workloads since making it an explicit go-to-market motion. Migration of sophisticated workloads is a combination of factors, including dissatisfaction with an incumbent and the attractiveness of a new cloud provider like **DigitalOcean**. As we build out our cloud capabilities, such as advanced networking, various **Droplet configurations**, storage like cold storage, and auto-scaling of our **DBaaS**, we become more attractive for migrations. While a single incident doesn't necessarily cause major shifts, these factors, along with enterprise-ready capabilities, make us more attractive for incoming migration. **AI-native workloads** are typically new, while many of the cloud workloads are migrations from other hyperscaler clouds.

Capacity Planning and Data Center Expansion

- **Q: Do you have the capacity to serve some of the eight-figure contracts signed after the quarter close in the coming weeks, or is the 30-megawatts of new data center capacity geared toward that, and can you provide context for the go-live times for these big contracts and how to think about the shape of 2026?**
- A: Some of these customers are already doing business with us. We are bringing some capacity online in our existing data center. The visibility with these customers and their inference scale-up is why we've expanded data center capacity. These data centers will come online progressively through **2026**. We have a build schedule and work closely with providers to bring this capacity online, but we don't have to wait for new capacity to light up these workloads. This visibility into the inference adoption of our **AI-native customers** is the reason for expanding our data center footprint. Most of the capacity will come online in the first half of next year. We'll be paying some of the **NRCs** for these build-outs in the fourth quarter, with capacity becoming available in the months and quarters following that. There's incremental time to deploy the **GPUs** and for customers to ramp, so we expect the revenue ramp to be relatively smooth over the year, with a fair bit of capacity coming online in the first half.

Competitive Strategy and Product Development

- **Q: A larger peer has acquired cloud capabilities, including a Python notebook and reinforcement learning for agents. What's your take on that, does it give credence to your strategy, and how do you see competition evolving as you continue to cater towards customers on the market?**
- A: Our strategy starts with a deep understanding of our customers and what it takes to serve them well, which has deepened over the last six months. We've had the **Python Notebook** capability for a couple of years. Our strategy targets **AI-native companies** building real businesses and running models in inferencing mode, requiring agentic workflow capabilities, storage, databases, authentication, authorization, and orchestration from a **Kubernetes** perspective. They need a unified agentic cloud stack, which we provide. We've been executing on our strategy, and **slide 7** shows the richness of our stack from infrastructure to middleware. **Slide 12** shows how much we've enhanced that since the last earnings call, with features lit up in a matter of days based on customer feedback. Building on our strength in software differentiation and leveraging our 12-year-old full-stack, general-purpose cloud puts us in a favorable position for scaling **AI-native companies**. Hardware is a part of it, but as companies become more sophisticated, the center of gravity will shift from hardware and networking to the software stack. We feel good about where we are and will be aggressive in adding new functionality.
- **Q: Are there any other factors that you guys take into account when deciding what to build next, like competitive positioning, or how do you balance the idea of what a customer wants with competitive positioning and long-term revenue opportunity, especially in something that could be seen as more experimental?**
- A: We are competitor aware, but customer-obsessed. That strategy has worked for us, especially in the fast-evolving **AI** landscape, where we've been disciplined in not chasing training workloads. We see an opportunity to decisively move to take a pole position in the world of inferencing and offer a software platform that combines the raw power of having the best of breed flexible **AI** infrastructure and combine it with where the puck is going, which is there is a whole generation, like **20 years** of app developed - applications that have been developed over the last **20 years** will need to be replaced and modernized with agents and agentic workflows. Managing that whole lifecycle is creating problems for companies to build, operate, and manage. There's a phenomenal opportunity for us to be one of the first movers into the world of agent development lifecycle, and that is exactly what we are focused on. While it's important to be aware of what our competition is doing, I feel very confident that we have unmatched software expertise and depth of our platform, as you can see from **slide 7** and **12**. If we keep doing what we're doing, our strategy is resonating with the **AI-native** and these are customers that are doubling and tripling their footprint on a month by month basis. If we can keep pace with them and keep shipping at their speed, I think everything else is going to take care of itself.

Net Dollar Retention (NDR) and AI Revenue

- **Q: AI is not included in net dollar retention, but are you thinking about starting to include it as it becomes a larger part of the business and more predictable?**
- A: We are looking very hard at how to incorporate the resilient growth of inferencing into our metrics. **NDR** is just - is one metric and it captures - it's used a lot in the **SaaS** environment. It captures the ongoing growth of a customer. To date, we haven't included it in **NDR** because a lot of the early traction that we were getting and I think and a lot of folks were getting in the industry, was more project-based and more experimental. As we're seeing customers like some of the ones that **Paddy** mentioned, like **Paul** and others, they are bringing scaled workloads to us where they have more predictability in their demand and growth. We believe it's appropriate to start to figure out how to include that. We'll likely revisit this once we get into the beginning of next year and we can have a better sense for the **2026** outlook. I mean, we're providing more specific guidance, but what would I say, the key takeaway is that **AI revenue** that we're seeing now that we're getting committed contracts for is behaving more like the traditional cloud where customers come in with scaled production workloads and then they have more visibility into the growth of that workload over time. Hence a metric like **NDR** becomes more relevant and we're confident that in doing that, that would be additive to our communications of the resilience of the growth that we're seeing.
- **Q: What other puts and takes are you considering as you look to drive NDR back over 100%?**
- A: Expansion is definitely the driver of the growth. If you look at the big customers, the customers that spend **\$100,000** or more in **ARR**, and as we showed, it gets better even as you get bigger spenders, they're driving a lot of our growth, and as you would expect, the **NDR** is better. The thing that people lose sight of, I think when they think about the **DigitalOcean** business is they forget that we have **640,000** plus customers and **450,000** or so or more thousand of those are effectively a paid premium, they're small customers, they spend **\$10, \$15** a month and many of them stay on for long periods of time. The average age of that cohort is something like **4.5 years**, but you also have a lot of customers come and go, they experiment. I mean, so the **NDR** of that paid premium segment of our customers is below **\$100**. And so, that weights down the overall **NDR** of the company and it masks the fact that with our largest customers we're actually seeing very, very strong growth, driven by increased expansion. That's another metric that we're thinking about is because we're blending this, where the real growth engine is for the company, we're blending that **NDR** with the **NDR** of a giant cohort, which is fantastic for us because it's - it gives us access to a lot of developers, but it's just by its nature it's going to have like slightly below **100 NDR**. We think that's masking a lot of the underlying performance that we're seeing.
- **Q: Given the strong growth trajectory and the confidence, can you just talk about where you think the net of both sort of explicit CapEx plus your equipment leasing, what the sum total of that could be as you look over the next couple of years in dollar terms?**
- A: As we said, we're trying to give preliminary guidance for **2026** to give the directional kind of growth rates. And we feel very confident that we can deliver that **18%, 20% revenue growth** while still delivering the kind of the mid-to-high teens in adjusted free cash flow. And when I say that number, that's the levered number that I'm referring to. And I think at this point that the market is evolving so fast that it's hard to say what the **CapEx** would be or what the impact on adjusted free cash flow margin would be even in the second half of next year, much less **2027** or beyond. What I can tell you is we - and you've seen us in terms of our behavior, we didn't chase the training opportunity. We didn't pursue what we viewed as revenue that we weren't sure how durable that was going to be for us, and we didn't know if we had a competitive differentiation there. But what we said is where we see opportunities to deliver durable revenue growth with a differentiated product that has good returns, that we'll make investments to pursue that. And you've seen that with our willingness to take down additional data center capacity and secure new **GPUs**. So, I'd say what you can expect is continued disciplined behavior where we're trying to drive durable revenue growth, while maintaining attractive free cash flow margin.

Q: How large could you see that delta growing between adjusted free cash flow and your adjusted unlevered free cash flow margins over the next few years?

- A: As we said, we're trying to give preliminary guidance for **2026** to give the directional kind of growth rates. And we feel very confident that we can deliver that **18%, 20% revenue growth** while still delivering the kind of the mid-to-high teens in adjusted free cash flow. And when I say that number, that's the levered number that I'm referring to. And I think at this point that the market is evolving so fast that it's hard to say what the **CapEx** would be or what the impact on adjusted free cash flow margin would be even in the second half of next year, much less **2027** or beyond. What I can tell you is we - and you've seen us in terms of our behavior, we didn't chase the training opportunity. We didn't pursue what we viewed as revenue that we weren't sure how durable that was going to be for us, and we didn't know if we had a competitive differentiation there. But what we said is where we see opportunities to deliver durable revenue growth with a differentiated product that has good returns, that we'll make investments to pursue that. And you've seen that with our willingness to take down additional data center capacity and secure new **GPUs**. So, I'd say what you can expect is continued disciplined behavior where we're trying to drive durable revenue growth, while maintaining attractive free cash flow margin.

Leasing and Gross Margins

- **Q: I want to ask about the transition to leveraging leasing and how the gross margins of data centers and equipment you own and operate compared to gross margins from leased capacity?**
- A: Just to clarify there, we don't - we lease all of our data centers. We're a co-location tenant in each of our data centers. We don't own any. So, the taking down of additional data center capacity that we've just referred to will behave the same way that it did when we took down the **Atlanta** data center earlier this year, and when we took down the **Sydney** data center several years ago. And what that does is our costs, if you think about our cost of

goods, are variable with revenue over the long-term, but in the very short term, they're somewhat lumpy. You take down an incremental data center, you have, not only do you have incremental upfront costs and **NRC** that happens before you're generating revenue, but the day you turn on the data center,

Show The Company's Guidance And Outlook

DigitalOcean's Q3 2025 earnings call provided financial guidance for Q4 2025 and the full year 2025, as well as an outlook for 2026. The company expects revenue growth, strong adjusted EBITDA margins, and non-GAAP diluted earnings per share. They are also making investments in data centers and GPU capacity to accelerate growth.

Quarterly Guidance

- Revenue is expected to be in the range of \$237 million to \$238 million , which is approximately 16% year-over-year growth.
- Adjusted EBITDA margins are expected to be in the range of 38.5% to 39.5%.
- Non-GAAP diluted earnings per share are expected to be \$0.35 to \$0.40 , based on approximately 111 million to 112 million in weighted average fully diluted shares outstanding.
- The non-GAAP diluted earnings per share guidance includes a projected impact of a range of about \$0.05 to \$0.10 reduction from the net impact of refinancing actions in Q3.

Annual Guidance

- Revenue is projected to be \$896 million to \$897 million , representing approximately 15% year-over-year growth.
- Adjusted EBITDA margin is expected to be approximately 41%.
- Non-GAAP diluted earnings per share are expected to be \$2.00 to \$2.05 , based on approximately 106 million to 107 million in weighted average fully diluted shares outstanding.
- The non-GAAP diluted earnings per share guidance includes a projected impact of a range of about \$0.15 to \$0.20 reduction for the full year from the net impact of the Q3 refinancing actions.
- Adjusted free cash flow margin of 18% to 19% is projected.

Outlook

- The company expects to comfortably deliver 18% to 20% growth in 2026 , achieving their 2027 revenue growth targets a full year earlier than projected.
- Adjusted EBITDA margins are anticipated to be in the high-30s to 40% .
- Adjusted free cash flow is expected to be in the mid-to-high teens .
- Net leverage is anticipated to end 2026 in the mid 3s range , including the impact on net debt from any incremental lease up.

Event Details

Date: 2025-11-05
Company: DigitalOcean Holdings, Inc.
Ticker: DOCN-US

Company Participants

Melanie Strate - DigitalOcean Holdings, Inc., Head-Investor Relations
Padmanabhan T. Srinivasan - DigitalOcean Holdings, Inc., Chief Executive Officer & Director
W. Matthew Steinfort - DigitalOcean Holdings, Inc., Chief Financial Officer

Other Participants

Gabriela Borges - Analyst
Radi Sultan - Analyst
Josh Baer - Analyst
Kingsley Crane - Analyst
James E. Fish - Analyst
Matthew Quigley - Analyst
Mark Zhang - Analyst
Wamsi Mohan - Analyst
Robert Galvin - Analyst

MANAGEMENT DISCUSSION SECTION

Operator

Thank you for standing by. My name is Rebecca, and I will be your conference operator today. At this time, I would like to welcome everyone to the DigitalOcean Third Quarter 2025 Earnings Conference Call. All lines have been placed on mute to prevent any background noise. After the speakers' remarks, there will be a question-and-answer session.

Thank you. I would like to turn the call over to Melanie Strate, Head of Investor Relations. Please go ahead.

Melanie Strate

Thank you, Rebecca, and good morning. Thank you all for joining us today to review DigitalOcean third quarter 2025 financial results. Joining me on the call today are Paddy Srinivasan, our Chief Executive Officer and Matt Steinfort, our Chief Financial Officer.

Before we begin, let me remind you that certain statements made on the call today may be considered forward-looking statements, which reflect management's best judgment based on currently available information. Our actual results may differ materially from those projected in these forward-looking statements, including our financial outlook. I direct your attention to the risk factors contained in our filings with the SEC, as well as those referenced in today's press release that is posted on our website. DigitalOcean expressly disclaims any obligation or undertaking to release publicly any updates or revisions to any forward-looking statements made today.

Additionally, non-GAAP financial measures will be discussed on this conference call. And reconciliations to the most directly comparable GAAP financial measures can be found in today's earnings press release as well as in our investor presentation that outlines the financial discussion on today's call. A webcast of today's call is also available on the IR section of our website.

And with that, I will turn the call over to Paddy.

Padmanabhan T. Srinivasan

Thank you, Melanie. Good morning, everyone, and thank you for joining us today, as we review our third quarter results. I'm very excited to share our results for the quarter and to give you an update on the progress that we are making against the goals that we articulated earlier this year during our April Investor Day. **Our performance this quarter was very strong, we exceeded our Q3 guidance on both revenue and profitability metrics**, delivering 16% revenue growth and the highest incremental organic ARR in the company's history, while generating 21% trailing 12-month adjusted free cash flow margins.

We continued innovation in our comprehensive agentic cloud to support the needs of scaling AI and digital native enterprise customers, making sure there is no reason our higher spending customers ever need to leave our platform. We augmented our industry-leading product-led growth engine with focused direct sales motion, driving customers to migrate workloads from the hyperscalers to our platform, and building traction with direct AI-native customers. **This progress** is evident in the rapid growth of our largest customers, and their increasing willingness to sign committed contracts with us. With **customers having more than \$1 million in annualized run rate, reaching \$110 million in ARR, growing 72% year-over-year** and with multiple customers signing eight figure committed contracts after the quarter close.

The demand for our agentic cloud has exceeded our supply. Our performance and the visibility we have into demand gives us the confidence both to increase our 2025 and 2026 revenue and adjusted free cash flow outlook, and to also increase our investment in data centers and GPU capacity to further accelerate growth, while maintaining attractive margin.

I will now dive deeper into all of this, starting with our third quarter financial results as highlighted on slide 10 of our earnings deck. **Q3 revenue hit \$230 million, up 16% year-over-year**, marking the highest growth since Q3 2023. We delivered our highest organic incremental ARR in company's history at \$44 million. This

growth was driven by a balanced performance across our comprehensive agentic cloud platform as direct AI revenue more than doubled year-over-year for the fifth consecutive quarter, and our general-purpose cloud products saw the highest incremental organic ARR since Q2 of 2022.

We delivered this accelerating revenue growth in Q3, while exceeding our profitability guidance and materially strengthening our balance sheet. Adjusted EBITDA and non-GAAP earnings per share were both well above guidance on the back of strong execution and **we delivered a strong 21%** trailing-12 month adjusted free cash flow margin as we introduced equipment leasing into our financial toolkit in Q3, to better align the timing of our investments with our revenue. To give us further flexibility to invest in growth, we also repurchased the majority of our 2026 convert in the quarter, strengthening our balance sheet.

The primary drivers behind our accelerating top line growth are threefold; number one, the increasing momentum we are seeing with AI-native customers; next, the material traction we continue to generate with our highest spend digital native enterprise customers; and finally, the continued strength we are seeing in revenue from new customers.

Our Unified Gradient AI agentic cloud, **which is outlined on slide 7 of our investor presentation, is getting increasing traction with larger, well-funded AI-native companies** that are in inference mode. These scaling companies increasingly leverage our unified agentic cloud with many of our top customers already leveraging both AI and general-purpose cloud capabilities, and with many more having at least starting to test and experiment with AI on our platform.

Evidence of this traction is in the growth rates of our highest spending customers. Revenue from these customers who are at \$100,000 plus **annual run rate grew 41% year-over-year, increasing to 26% of total revenue. Growth is even higher for our largest digital-native enterprise customers**, as the more our customers are spending, the faster they're growing on DO (00:07:09). The charts on slide 11 show that **our customers with greater than \$500,000 and greater than \$1 million in annualized run rate grew revenue 55% and 72%, respectively**, providing clear evidence that our increasing ability to not just attract, but also retain and grow our largest customers demonstrating that customers can keep scaling on our platform and never have a reason to leave.

Let me now dive deeper into this traction using slide 12 as the backdrop to illustrate just how much progress we have made since the last earnings call. I will start with our AI infrastructure on the bottom right, which is a full-stack inference platform targeting AI-native customer that have their own models that they want to tune, optimize and run in inference mode. These customers select our platform for our full set of capabilities where we combine a powerful lineup of GPUs **that are available in both bare metal and droplet configurations, including inference-optimized droplet with advanced inference performance optimization like paged attention, FlashAttention, FP8 quantization, Speculative Decoding, model operations management, reduced** time for first token and compelling TCO economics.

Our AI infrastructure provides comprehensive hardware plus software infrastructure for AI-native companies that are scaling up real world inference workloads globally on DO. Fal.ai or fal, a generative media model platform that provides text to image and text to video models for major customers such as Canva, Shopify, Perplexity and More is a great example of a customer **that is taking advantage of our unified agentic cloud**. They leverage a range of our AI infrastructure solutions, including GPU droplets, both to host their media models in production, serving their end customers, and to do research and fine-tune.

Fal is more than just an important customer, as we have come together in a strategic partnership to accelerate generative AI content creation by making image and audio generation more accessible to startups and enterprises. Through this partnership, fal will host and run hundreds of its models on DigitalOcean's

infrastructure, powering applications across creative and enterprise use cases. This means customers can create agents that understand and generate not only text, but also images, data, and other forms of input, significantly expanding the range of real world problems our customers can solve.

NewsBreak is another example of an AI-native customer leveraging our unified agentic cloud. Driving the next generation of digital media, NewsBreak delivers timely and relevant local news and information to 40 million monthly active users. NewsBreak's AI-powered infrastructure makes sophisticated personalization accessible to mainstream users nationwide. They utilize our AI infrastructure to train and deploy complex recommender system and natural language processing models that are foundational to their product. **Our AI infrastructure's high throughput and memory capacity** are critical for running inference at scale, which allows them to perform real-time content ranking and ad placement for millions of concurrent users.

Gradient AI agentic cloud unifies our integrated AI capabilities with our full-stack general-purpose cloud, which we have been optimizing for over a decade, enabling NewsBreak to pre-process their work on our CPU droplet and run their vector search service in advance of running their AI workloads, optimizing both cost and performance.

Network File Storage or NFS, which delivers high throughput performance for both GPU and non-GPU droplets, is an example of a unified agentic cloud capability. Customers can now attach and provision storage in just minutes, accelerating time to value by eliminating idle time. With seamless integration into our Kubernetes engine, NFS makes it easier than ever to scale applications and workloads, while maintaining speed, reliability and efficiency across environments.

Moving up the stack outlined on slide 7, the AI platform layer on the middle right is typically leveraged by companies that are users or consumers of AI that are looking to build agentic application without having to directly manage the infrastructure. As we know, the future of AI is an agent and agentic workflows, which is a natural evolutionary step for all staff and other applications. We continue to evolve our AI platform as the foundation for building and deploying these intelligent agents and powering complex enterprise agentic workflows. It now supports serverless inferencing across the most popular model, including OpenAI, Anthropic, Mistral, LLama, DeepSeek and others, including new generative media models from fal.

We've added a powerful knowledge-based service that lets customers bring their own data and improve accuracy, along with built-in guardrails for safety, visual agent orchestration and enterprise-grade features like observability, Git integration and auto scaling. Together, these capabilities make our Gradient AI agentic cloud platform one of the most intuitive and complete platforms for taking AI agents from prototype to production. These key capabilities help companies develop and operate AI agents fleet (00:13:40) and manage their full lifecycle of these agents seamlessly from a single platform, while leveraging the best of breed AI models from various providers.

We are particularly excited about a major customer we signed for our AI platform after the Q3 quarter close. This customer is a global digital systems integrator, who signed an eight figure per year multiyear contract to leverage our agentic cloud to drive AI transformation for its digital-native enterprise customer base with a specific focus on agentifying the full software engineering lifecycle, including planning, backlog, and roadmap management, release planning, release execution, and customer support. **We'll provide more information on this exciting customer** after we formally announce the partnership in the upcoming days.

The AI platform layer continues to also gain broader momentum, with over 19,000 agents created so far, of which more than 7,000 are already in production. **One specific customer, SHAKAZAMBA (00:14:52), an Italian leader in GDPR compliant**, ethical and secure AI solutions across Europe, chose to leverage the Gradient AI agentic cloud over the hyperscalers. By using our platform, they're now able to create and roll out

agents to automate customer support, knowledge management, and content creation while reducing development time and costs associated with the agent lifecycle.

This quarter, **we also expanded our AI ecosystem with the launch of the DigitalOcean AI Partner Program** with several of our partners outlined on Slide 13. This is a major step in empowering AI and digital-native enterprises **that are building and scaling their businesses**, leveraging AI. These companies don't have time for a fragmented infrastructure, they instead want a unified cloud and an AI platform that lets them seamlessly build and scale intelligent applications using agents. This new partner program brings together AI-native companies, integrators, and the venture ecosystem to help these builders reach more customers, accelerate innovation and amplify their global reach. Combined with our AI platform and infrastructure, this ecosystem makes DigitalOcean the go-to-destination for these AI-native businesses who want simplicity, scalability, and reach without the hyperscale complexity.

In Q3, **we continued to deliver product innovation in our core cloud-stack** to support our highest spending customers by meeting their needs, as they scale their business on DO. One such example of a digital-native enterprise customer scaling rapidly on DO is Bright Data, a leading provider of web datasets to global frontier LLM labs for training AI models. Bright Data leverages various components of our agentic cloud to scale high-volume global workloads on our platform.

VPN Super, **who develops trusted VPN and Security Solution, is the most downloaded VPN app in the world, is another digital-native enterprise growing on our platform.** VPN Super empowers millions of users across the globe to browse securely and privately regardless of their location. They signed a seven figure deal to migrate multiple workloads to this DigitalOcean and their selected deal for our ability to handle large traffic spikes, platform reliability and our global scale. **These growing customer require general purpose cloud capabilities** that grow with (00:17:42) their business and we delivered a number of these new features during the quarter, as you can see, highlighted on slide 12 of our earnings presentation.

For example, we recently introduced Spaces Cold Storage, an enterprise-grade object storage solution designed for customers managing data at massive scale, with support for hundreds of petabytes and billions of object per bucket. It offers free retrieval, predictable low cost and immediate access to data, eliminating the trade-off between affordability and performance. This cold storage is secure, reliable and resilient, providing our customers with the confidence to store and access mission-critical datasets seamlessly as their needs grow.

During the quarter, **we also enhanced our managed databases offering with automated storage auto scaling**, enabling customers to scale seamlessly as their data needs grow. When capacity thresholds are reached, storage automatically scale in 10 gigabytes increments or higher, with zero downtime and no disruption to workload. This feature is available across all major database engines, including MongoDB, PostgreSQL, MySQL and is fully customizable, allowing customers to set thresholds starting at 20% utilization. With a simple pay-as-you-go model, auto scaling eliminates the burden of manual intervention, ensuring that application scale reliably and cost effectively.

The steady stream of new features is resonating with our AI and digital-native enterprise customers. Over 35% of our customers with more than \$100,000 in ARR have adopted at least one of our new features released over the past year. **And those customers having adopted at least one of these new products, have seen a several hundred basis points increase in their growth rate after adopting the new product.**

Our strong performance, our growing momentum through the first three quarters and the visibility that we now have into demand gives us the confidence to raise our near and medium-term growth outlook. **We are raising our full-year 2025 guidance on both revenue and margin**, and we now expect to achieve our 18%

to 20% 2027 revenue growth target in 2026, a full year earlier than we had projected. It has also given us the confidence to accelerate our investments to drive growth in 2026 and beyond.

When we outlined our 2027 growth objectives this past April, we indicated that we would increase our investments as we saw opportunities to accelerate our growth. **We're now seeing more demand than we can support with our existing capacity, which is evident by us** having signed multiple eight figure committed contracts after the quarter ended. That will materially increase our RPO in Q4. With this increased conviction, we began to put the foundational elements in place in Q3 to even further accelerate our growth. **We started ordering more GPU capacity to meet the growing inference demand** we are seeing from our AI-native customers. **We also secured around 30 megawatts of incremental data center capacity to support growth in 2026** and beyond. We added equipment financing to better align our investments with revenue. We ramped engineering resources to accelerate our unified agentic cloud roadmap and continued our targeted investment in new sales and marketing initiatives to complement our industry-leading product-led growth engine.

These investments will build on the success we have seen today and will set us up for a strong 2026 and 2027. Our Q4 and 2025 full-year guidance implies a 16% exit 2025 growth rate. And while we won't provide 2026 guidance until our February earnings call, we expect to comfortably deliver 18% to 20% growth in 2026, achieving our 2027 growth target a full year earlier than previously projected. **We will deliver this growth while maintaining strong adjusted free cash flow margins in the mid-to-high teens.** Matt, will provide further color on these investments and the projected impact on our growth and profitability in his remarks.

As I said in my opening, **we delivered strong performance in Q3**, beating our guidance on both revenue and profitability. We are seeing momentum with our unified agentic-cloud and this momentum is evident in the rapid growth of our highest spending customers and demand is exceeding our current capacity. All of this gives us the conviction both to raise our 2025 and 2026 revenue and adjusted free cash flow outlook and to increase our investments to take advantage of the opportunity in front of us. We look forward to sharing more on our progress and our outlook for 2026 over the upcoming months.

Thank you. And I'll now turn it over to Matt.

W. Matthew Steinfort

Thanks, Paddy. Good morning, everyone, and thanks for joining us today. As Paddy discussed, **we are excited about our strong Q3 2025 performance.** We are gaining traction with our unified agentic cloud, **which is resulting in strong revenue growth from our highest spending customers.** And **we are seeing more demand** and satisfied with our current capacity. This momentum and visibility gives us conviction both to increase our 2025 and 2026 revenue and adjusted free cash flow outlooks, and to put in place the foundations to further accelerate growth in 2026 and beyond.

In my comments, I'll walk through our Q3 results in detail, share our fourth quarter and updated full year financial outlook, and also provide an update on our 2026 expectations. Starting with the top line, **revenue in the third quarter was \$230 million, up 16% year-over-year**, the highest revenue growth since Q3 of 2023. This growth was balanced across our unified agentic cloud and was primarily driven by increasing traction with our higher spending AI digital-native enterprise customers, with steady contributions from our product-led growth engine.

We continued to see strong AI/ML revenue growth in Q3, with AI revenue more than doubling year-over-year, which it has done every quarter since we launched our AI platform. We also delivered the highest incremental organic ARR in company history at \$44 million, bringing ARR to \$919 million. With rapid product

innovation across our unified agentic cloud platform and with the strategic go-to-market investments we made this year proving to be effective, **we're having increasing success**, attracting and growing larger, well-funded AI and digital-led enterprise customers. This is evident in the revenue from our customers, whose annualized run rate revenue (00:25:07) quarter was greater than \$100,000, which now represents 26% of overall revenue, growing 41% year-over-year, 200 basis points higher than the growth we saw from that cohort in the prior year. Adding to this growth, revenue from general-purpose cloud customers in their first 12 months on our platform continues to be strong and we have stabilized NDR, as **net dollar retention remained at 99% in the quarter, up 200 basis points from 97% in the third quarter of 2024.**

Turning to the P&L. While we accelerated revenue, we also delivered strong performance on all of our key profitability metrics. **Gross profit was \$137 million, up 17% year-over-year, with a 60% gross margin for the third quarter**, a 100 basis points higher than the prior year. **Adjusted EBITDA was \$100 million, a 15% increase** year-over-year, and an adjusted EBITDA margin of 43%. Non-**GAAP diluted net income per share was \$0.54, a 4% increase year-over-year.** **This result was impacted by the \$625 million convertible note** we issued in August, the repurchase of \$1.19 billion of our 2026 notes, and the interest expense from the \$380 million drawn on the term loan A component of our existing credit facility.

The net impact on non-GAAP net income per share of these balance sheet activities was a reduction of \$0.05 in Q3. And excluding these charges, non-GAAP diluted net income per share would have been \$0.59. **GAAP diluted net income per share was \$1.51, a 358% increase year-over-year.** **This increase** is primarily driven by the one-time reversal of our tax valuation allowance and gain on debt extinguishment, which is slightly offset by the impact of our new debt structure.

Q3 adjusted free cash flow was \$85 million or 37% of revenue, **which is up significantly from the prior year's** \$19 million or 10% of growth, and on a trailing 12-month basis was 21% of revenue. **This increase** was driven in part by the equipment financing in Q3. During the quarter, we entered into an equipment financing arrangement with the third-party financial institution for \$28 million to better align our investments with the future revenue that they will generate. Absent the leasing of this equipment, our adjusted free cash flow would have been 25% of revenue in Q3.

Turning to the balance sheet, **we strengthened our balance sheet by repurchasing approximately 80% of our 2026 convertible notes through a combination of the issuance of a new \$625 million 2030 convertible note offering**, a \$380 million drawdown on the term loan (00:27:49) component of our existing credit facility, and approximately \$230 million of cash. Following these actions, our cash and cash equivalents balance remained healthy at \$237 million, and the combination of cash on hand, remaining term loan A capacity and projected cash flow generation is collectively more than the remaining balance of our outstanding 2026 convertible notes.

We repurchased \$2.9 million of shares in Q3, buying back approximately 101,000 shares, bringing our cumulative share repurchase since IPO to \$1.6 billion and 34.9 million shares through September 30, 2025. These Q3 repurchases completed our 2024 buyback program, and we will operate our repurchase program through July 31, 2027, under the new \$100 million authorization we announced during the quarter. During the quarter, we repurchased a portion of our 0% coupon 2026 convertible notes in part with an interest bearing term loan A that is initially at SOFR plus 175 basis points for roughly 6.1%. As a result, we now project to have modest interest expense in the near to medium term, where interest expense was previously immaterial.

Given this, we have added a new disclosure metric for unlevered adjusted free cash flow, which we will provide in addition to the current adjusted free cash flow metric, which is a levered adjusted free cash flow. We believe that unlevered adjusted free cash flow is an important metric, as it provides a clear view of our

cash generation before the impact of financing decisions. And many investors and analysts use this unlevered, adjusted free cash flow as the basis for their enterprise value calculation.

Our Q3 unlevered adjusted free cash flow was \$85 million or 35% – 37% of revenue. **The strong demand** we've seen across our unified agentic cloud and the traction we are seeing with our higher spending AI and digital-native enterprise coupled with the increased visibility we have from having signed multiple eight figure committed contracts after Q3 close gives us the confidence to raise our outlook on both revenue and adjusted free cash flow margin for both 2025 and 2026.

For the fourth quarter of 2025, we expect revenue to be in the range of \$237 million to \$238 million, **which is approximately 16% year-over-year growth**. For the full year 2025, **we project revenue of \$896 million to \$897 million, representing approximately 15% year-over-year growth, an incremental 100 basis points higher than our prior guide**. For the fourth quarter of 2025, we expect our adjusted EBITDA margins to be in the range of 38.5% to 39.5% with an adjusted EBITDA margin of approximately 41% for the full year.

For the fourth quarter of 2025, we expect non-GAAP diluted earnings per share to be \$0.35 to \$0.40 based on approximately 111 million to 112 million in weighted average fully diluted shares outstanding. For the full year 2025, we expect non-GAAP diluted earnings per share to be \$2.00 to \$2.05 based on approximately 106 million to 107 million in weighted average fully diluted shares outstanding. For Q4 and full year, **non-GAAP diluted earnings per share guidance included a projected impact of a range of about \$0.05 to \$0.10 reduction in Q4**, and about \$0.15 to \$0.20 reduction for the full year from the net impact of our Q3 refinancing actions. We project a full year adjusted free cash flow margin of 18% to 19%.

Looking further ahead, I would also like to provide a brief update on our 2026 outlook. While we will provide more fulsome details on 2026 expectations during our earnings call in February, we have already begun to put the foundations in place to further accelerate growth. Given **our momentum and the increased visibility into demand on the back of several recent customer wins**, we have committed investment in additional data centers and GPU capacity that will come online over the course of 2026. **That will accelerate growth** ahead of our previously communicated timeline. We have signed leases for approximately 30 megawatts of incremental data center capacity across several new data centers that will commence over the course of 2026. These new data centers and our corresponding investments in incremental GPU capacity will enable us to comfortably deliver 18% to 20% growth in 2026, achieving our 2027 revenue growth targets a full year earlier than we had projected.

And while **our COGS and operating expenses will increase in early-2026**, as we ramp into our new data center capacity, we anticipate delivering high-30s to 40% adjusted EBITDA margins, while maintaining mid to high teens adjusted free cash flow. **We also remain committed to maintaining a healthy balance sheet** and we anticipate that our net leverage will end 2026 in the mid 3s range (00:32:55), including the impact on net debt from any incremental lease up.

We look forward to sharing more on the traction we are getting with our unified agentic cloud, the growth we are seeing from our highest spending customers, the investments we are making to further accelerate growth and our outlook for 2026 and beyond when we get together again in February.

That concludes our prepared remarks and we'll now open the call for Q&A.

QUESTION AND ANSWER SECTION

Operator

. Your first question comes from the line of Gabriela Borges with Goldman Sachs.

Analyst:Gabriela Borges

Question – Gabriela Borges: Hey, good morning. Thank you for taking the question and congratulations on a really exciting 2026 preliminary forecast. Paddy and Matt, I want to ask you about the multiple eight figure committed contracts that you're talking about. Tell us a little bit about this cohort of customers? To what extent does it overlap with some of the AI revenue that you're talking about? I know you've been working with the private equity community as well, you've talked about migration. So, maybe just a little bit about the type of customer that's signing the eight figure contract and the extent to which? I know in the past the AI cohort has been much more flecky (00:34:25) in its ability to ramp up and down. And so, I'm trying to understand the intersection between those two cohorts. Thank you.

Answer – Padmanabhan T. Srinivasan: Yeah, thank you, Gabriela. Great questions. So, the eight figure commitment contracts that we just talked about come in different forms. Primarily, these are AI-native companies that are looking to take advantage of our infrastructure, as well as the other customer that I was just talking about for our AI platform, looking to build a series of agentic experiences for software engineering, taking advantage of our Gradient AI platform layer.

So, it's a combination of all of these. And as I was explaining in my prepared remarks, it is increasingly getting difficult to just separate out where AI starts and stops and where core cloud begins, because most of the customers that are starting their experience with DigitalOcean from the AI side are increasingly using our various storage artifacts like Network File System or the VPC capabilities or a number of the networking capabilities and things like that.

So, the crossover is becoming more and more between the AI side of our platform and cloud. So, that's why we are now starting to see a more unified cloud platform from us, which we are calling as the agentic cloud. So, a lot of these commitments and contracts that we are starting to take on now typically start with AI, but also spill over to our AI cloud side as well. So, what is exciting for us is that some of these customers start their journey with DO, using a fairly small proof-of-concept type of footprint, and now they're starting to scale. And this is also one of the many reasons why we are expanding our data center footprint so that **we can keep scaling with these customers.**

And the other attribute I want to call out here is that these AI workloads are predominantly inferencing, if not all, on the inferencing side. So, it is durable, it is predictable. And **we also have a great opportunity to keep scaling with these customers** as they find real-world traction and scale globally. So, that's why it is really important for us to start looking at our capacity as we place bets on some of these real marquee AI-native companies **that are finding traction with real end customers**, both on the consumer side as well as on the enterprise side.

Operator

Your next question comes from the line of Radi Sultan with UBS.

Analyst:Radi Sultan

Question – Radi Sultan: Awesome, yeah, thanks for taking the question, guys, and good to see the platform traction coming in ahead of schedule. I guess, for me, just AWS and Azure, both had some pretty high profile outages recently. And I'm just curious, like, is that having any near-term impact or catalyzing more migrations from the hyperscalers, that's driving more traction for Partner Network Connect or some of the other multi-cloud offerings? And then just curious, how many of those eight figure deals were migrations from the hyperscalers? Thank you.

Answer – Padmanabhan T. Srinivasan: Yeah. Thank you Radi for the question. So, **we've been seeing a steady increase in migration workloads** since we made it into an explicit go-to-market motion. And as you know, migration of sophisticated workloads, it's always a combination of factors, right? It's seldom, oh, we see a disruption from a cloud provider and we're just going to move a fairly sophisticated global workload, but it is a combination of factors. **Some are driven by dissatisfaction with an incumbent**, but mostly it is driven by something they find attractive in a new cloud provider like DigitalOcean.

So, as we have started building out our cloud capabilities, especially the ones that I described in slide 12 of our deck, like more advanced networking, various flavors of our Droplet configuration, our storage, like cold storage, is a really, really important capability that many of our large customers with sophisticated workloads have been asking us for, the auto scaling of our DBaaS. I mean, these are all very fundamentally building block type of capabilities for attracting more migration workloads, not to mention some of the stuff we did in the last couple of quarters, like virtual private cloud and Direct Connect and things like that.

So, even though a single incident doesn't necessarily precipitate major shifts in workloads, these are all paper cuts and us having these other digital-native enterprise-ready capabilities just makes us all the more attractive to incoming migration. And the AI-native workloads typically are new workloads that are starting on our platform, but many of the workloads that we are seeing that I talked about during my prepared remarks on the cloud are migrations coming from various other hyperscaler clouds.

Operator

Your next question comes the line of Josh Baer with Morgan Stanley.

Analyst: Josh Baer

Question – Josh Baer: Great. Thanks for the question and congrats on the acceleration. I wanted to follow up on the eight figure contracts signed after quarter close. Just I'm wondering, do you have capacity to serve some of those in the coming weeks or is that all really what the 30-megawatts of new data center capacity is geared toward? And then hoping to get a little sense of like the ramp in that capacity, basically any context for the go-live times for these big contracts and like how to think about the shape of 2026?

Answer – Padmanabhan T. Srinivasan: Yeah. So, I'll get started, Matt, and then you can chime in. So, from the ramp-up perspective, some of these customers are already doing business with us. And as we think about the capacity, there is some capacity we are bringing online in our existing data center. And then, of course, a lot of the visibility that we now have with these customers and their inference scale up is the reason why **we have taken up expanded data center capacity**. And these data centers will come online progressively through 2026, right. So, we do have a build schedule from these providers, and we work very, very closely with them to make sure that we get the warm shell scale and then we move in and we start racking our servers, and there's a lot of moving parts in terms of bringing this capacity online.

But we don't have to wait for this new capacity to start lighting up these workloads. As I said, we do have some capacity in our existing data center. So, it is a combination of these things. And you're absolutely right that this visibility into the inference adoption of our AI-native customers is the reason why **we are expanding our data center footprint.**

Answer – W. Matthew Steinfert: Yeah, and I'll just add that most of the capacity is going to come online in, call it (00:42:10), the first half of next year. In fact, you'll see in our fourth quarter financials and this is included in the guidance for the adjusted free cash flow we have for 2025 that we're going to be paying some of the NRCs for some of these build-outs in the fourth quarter. And so, we expect that the capacity will become online in the months and quarters following that. So, it'll be an early ramp of the data center capacity. But then clearly, you have incremental time to deploy the GPUs and for the customers to ramp. So, we expect **the revenue ramp to be to be relatively smooth over the course of the year**, but we'll be bringing on a fair bit of capacity in the first half.

Operator

Your next question comes from the line of Kingsley Crane with Canaccord Genuity.

Analyst:Kingsley Crane

Question – Kingsley Crane: Hi, thanks for taking my question, I want to echo my congrats. I'm sure (00:43:08) it's gratifying for the team. With a larger peer neocloud (00:43:12) peer has acquired a handful of cloud (00:43:14) capabilities over the past six months, including more recently a Python notebook, (00:43:19) reinforcement learning for agents. What's your take on that? Just to give credence to your strategy and how do you see competition evolving as you continue to cater towards customers on market?

Answer – Padmanabhan T. Srinivasan: Yeah, Kingsley, thank you for the question. So, we – our approach when we laid out our strategy in April, we start our strategy with a deep understanding of who our customers are and what it would take for us to serve them well. And of course, that understanding has deepened over the last six months as we have started working very, very closely with these customers. In terms of the Python Notebook, this was – this has been a capability that we have had for a couple of years now. And some of the storage enhancements that we are seeing in the market is also something that has been long, as part of our very rich and deep software stack.

So, if you take a step back and think about our strategy, our strategy is now, from an AI perspective, targeting AI-native companies that are building real businesses and running models in an inferencing mode, and these are real world applications that require not just GPU and inferencing capabilities, but they need agentic workflow capabilities. They need storage, databases, authentication, authorization, they need orchestration from a Kubernetes perspective. So, essentially, they need a unified agentic cloud stack, which is what we provide.

So, we have been executing on our strategy. And if you look at slide 7, you see the richness of the stack that we have built all the way from infrastructure on both cloud and AI to middleware with platform-as-a-service or the agentic development lifecycle. And slide 12 shows how much we have enhanced that since just the last earnings call. So, we have been super busy and every orange box you see on slide 12 has been a result of feedback that we're getting from customers real time, **some of these features have been lit up in just a matter of days**, and that is the power of really co-inventing some of these pieces, working hand-in-hand with our customers. **And I feel building on our strength, which is software differentiation and leveraging the**

strength of our 12-year old full-stack, **general purpose cloud really puts us in a very favorable position** when it comes to being attractive to these scaling AI-native companies.

Is hardware a part of it? Absolutely. But we think as companies become more and more sophisticated and they start serving real world enterprise needs, the center of gravity is going to shift from hardware and networking and move more and more towards the software stack. As we have seen in every wave that we have encountered over the last two or three decades. So, **we feel really good about where we are and we will be aggressive in adding new functionality into our platform as we start seeing opportunities** for those and from our customers and in the market.

Operator

Your next question comes from the line of Patrick Walravens with Citizens.

Unidentified speaker

Question – Unidentified speaker: Hi, thank you guys so much for taking the question and congrats on the quarter. This is Nick (00:47:00) on for Pat. Paddy, one for you. You kind of answered this already, but are there any other factors that you guys take into account when deciding what to build next? Like you mentioned that it's primarily customer driven, but are there like competitive positioning or, how do you balance the idea of what a customer wants with competitive positioning and long-term revenue opportunity, especially in something that could be seen as more experimental?

Answer – Padmanabhan T. Srinivasan: Yeah, Nick, thank you for the question. So, just for the avoidance of doubt, we are competitor aware, but customer obsessed. So, we – and that strategy has really worked for us, especially in the fast evolving AI landscape where we have been very, very disciplined in not chasing the bright, shiny training workloads and trying to be somebody that we are not. But we have been patient and now we see an opportunity to decisively move to take a pole position in the world of inferencing and offer a software platform that combines the raw power of having the best of breed flexible AI infrastructure and combine it with where the puck is going, which is there is a whole generation, like 20 years of app developed – applications that have been developed over the last 20 years will need to be replaced and modernized with agents and agentic workflows.

So, managing that whole lifecycle is starting to already create a tremendous amount of problems for companies to build, operate, and manage. So, we think there's a phenomenal opportunity for us to be one of the first movers into the world of agent development lifecycle, and that is exactly what we are focused on. So, while it is important to be aware of what our competition is doing, especially the neocloud (00:49:00), I feel very confident that we have unmatched software expertise and depth of our platform, as you can see from slide 7 and 12.

So, if we keep doing what we're doing, our strategy is resonating with the AI-native and these are customers **that are doubling and tripling their footprint on a month by month basis.** So, if we can keep pace with them and keep shipping at their speed, I think everything else is going to take care of itself.

Operator

Your next question comes from the line of James Fish with Piper Sandler.

Analyst:James E. Fish

Question – James E. Fish: Hey, guys, nice quarter. Just on the 2026 starting point, kudos on bringing that forward, understanding what's driving the confidence and visibility at this point. But how should we think about the parts underneath between core managed service and AI? And do we still expect that AI business to continue to double given you've done it for five straight quarters now? Thanks, guys.

Answer – W. Matthew Steinfort: Thanks, Fish. We think that if you look at the success that we're having, it's coming from a combination of things. **It's coming from our success with our largest customers, regardless of whether they're AI or core cloud, growing very, very rapidly with the million-dollar-plus customers growing 72%,** we think that continues. And a lot of the migration workloads that we've been working on and some of these longer term committed contracts are also on the core cloud side. So, **growth of our biggest customers** is kind of the lever number one.

Lever number two is, as you said, is AI growth. **It has been doubling every quarter** since we launched it and we're expecting that to continue. So, **we're going to get a big chunk of growth from AI** that will become a more material part of our business. It'll get into the mid-teens, maybe even high teens as a percent of revenue. And then **we're still generating very good incremental revenue from our product-led growth engine. Our customers in M1 to M12 are still at very high levels** relative to historical and it's really those three levers that give us confidence. And as you said, we have better visibility now than we've ever had. If you said how many eight figure committed contracts have we ever had in the company, it's – I don't know how many, it's not that many, it's maybe one or two. And now, we've got multiple that we've signed just in the last month or so. So, **we're very confident in the 18% to 20% revenue for 2026** and are excited to be able to pull that in a whole year.

Operator

Your next question comes from the line of Mike Cikos with Needham and Company.

Analyst:Matthew Quigley

Question – Matthew Quigley: Hey guys, this is Matt Quigley on for Mike Cikos over at Needham. Thanks for taking our questions and great to see **the strength of large deals**. I know AI is not included in net dollar retention, but are you thinking about starting to include it as it becomes a larger part of the business and more predictable? And what other puts and takes are you considering as you look to drive NDR back over 100%?

Answer – W. Matthew Steinfort: Yeah, it's a great question. We are looking very hard at how to incorporate the resilient growth of inferencing into our metrics. NDR is just – is one metric and it captures – it's used a lot in the SaaS environment. It captures the ongoing growth of a customer. To date, we haven't included it in NDR because a lot of the early traction that we were getting and I think and a lot of folks were getting in the industry, was more project-based and more experimental. As we're seeing customers like some of the ones that Paddy mentioned, like Paul and others, they are bringing scaled workloads to us where they have more predictability in their demand and growth. We believe it's appropriate to start to figure out how to include that.

We'll likely revisit this once we get into the beginning of next year and we can have a better sense for the 2026 outlook. I mean, we're providing more specific guidance, but what would I say, the key takeaway is that AI revenue that we're seeing now that we're getting committed contracts for is behaving more like the

traditional cloud where customers come in with scaled production workloads and then they have more visibility into the growth of that workload over time. Hence a metric like NDR becomes more relevant and we're confident that in doing that, that would be additive to our communications of the resilience of the growth that we're seeing.

Operator

Your next question comes from the line of Mark Zhang with Citigroup.

Analyst:Mark Zhang

Question – Mark Zhang: Hey, good morning, guys. Thanks for taking the questions. Maybe just on NDR, obviously (00:53:58) still at the 99%, but **we're seeing good expansion momentum and portfolio momentum**. So, can you maybe just walk through some of the key puts and takes there? And traditionally, I think your expansion activity has been masked (00:54:12) in that metric. So, can you maybe speak to some of the behaviors there and any discernible changes in customer behavior versus last quarter or 12 months ago, whether it's on pace of expansion or how they're expanding? Thank you.

Answer – Padmanabhan T. Srinivasan: Thanks, Mark. It's a great, great question. **Expansion is definitely the driver of the growth**. If you look at the big customers, the customers that spend \$100,000 or more in ARR, and as we showed, it gets better even as you get bigger spenders, **they're driving a lot of our growth**, and as you would expect, the NDR is better. The thing that people lose sight of, I think when they think about the DigitalOcean business is they forget that we have 640,000 plus customers and 450,000 or so or more thousand of those are effectively a paid premium, they're small customers, they spend \$10, \$15 a month and many of them stay on for long periods of time. The average age of that cohort is something like 4.5 years, but you also have a lot of customers come and go, they experiment.

I mean, so the NDR of that paid premium segment of our customers is below \$100. And so, that weights down the overall NDR of the company and it masks the fact that with our largest customers **we're actually seeing very, very strong growth, driven by increased expansion**. And that's another follow-up to the question prior to this. That's another metric that we're thinking about is because we're blending this, where the real growth engine is for the company, we're blending that NDR with the NDR of a giant cohort, which is fantastic for us because it's – it gives us access to a lot of developers, but it's just by its nature it's going to have like slightly below 100 NDR. We think **that's masking a lot of the underlying performance** that we're seeing.

Operator

Your next question comes from the line of Wamsi Mohan with Bank of America.

Analyst:Wamsi Mohan

Question – Wamsi Mohan: Yes, thank you so much. Nice results here. Just given **the strong growth trajectory and the confidence**. Can you just talk about where you think the net of both sort of explicit CapEx plus your equipment leasing, what the sum total of that could be as you look over the next couple of years in dollar terms? And how large could you see that delta growing between adjusted free cash flow and your adjusted unlevered free cash flow margins over the next few years? Thank you.

Answer – W. Matthew Steinfort: Hey, Wamsi. It's a great question. I think, as we said, we're trying to give preliminary guidance for 2026 to give the directional kind of growth rates. And we feel very confident that we can deliver that 18%, 20% revenue growth while still delivering the kind of the mid-to-high teens in adjusted free cash flow. And when I say that number, that's the levered number that I'm referring to. And I think at this point that the market is evolving so fast that it's hard to say what the CapEx would be or what the impact on adjusted free cash flow margin would be even in the second half of next year, much less 2027 or beyond.

What I can tell you is we – and you've seen us in terms of our behavior, we didn't chase the training opportunity. We didn't pursue what we viewed as revenue that we weren't sure how durable that was going to be for us, and we didn't know if we had a competitive differentiation there. But what we said is where we see opportunities to deliver durable revenue growth with a differentiated product **that has good returns**, that we'll make investments to pursue that. And you've seen that with our willingness to take down additional data center capacity and secure new GPUs. So, I'd say what you can expect is continued disciplined behavior where we're trying to drive durable revenue growth, while maintaining attractive free cash flow margin.

Operator

Your final question comes from the line of Robert Galvin with Stifel.

Analyst:Robert Galvin

Question – Robert Galvin: Hi, thanks for taking the question. I want to ask about the transition to leveraging leasing and how the gross margins of data centers and equipment you own and operate compared to gross margins from leased capacity? Thanks.

Answer – Padmanabhan T. Srinivasan: Right. So, just to clarify there, we don't – we lease all of our data centers. We're a co-location tenant in each of our data centers. We don't own any. So, the taking down of additional data center capacity that we've just referred to will behave the same way that it did when we took down the Atlanta data center earlier this year, and when we took down the Sydney data center several years ago. And what that does is our costs, if you think about our cost of goods, are variable with revenue over the long-term, but in the very short term, **they're somewhat lumpy**.

You take down an incremental data center, you have, not only do you have incremental upfront costs and NRC that happens before you're generating revenue, but the day you turn on the data center, you start paying for the space and some of the power. Then you build it out and kit it out with gear and you fill it up and it takes some period of time to generate the fully utilized revenue in that facility.

So, there's always **a lump of higher expenses in the beginning** when you turn on data center capacity and you grow into it and you grow into it over a series of even just a couple of quarters, and your gross margin gets back to what's more of a steady state gross margin. So, we expect that to happen in the beginning of next year of 2026. But we've factored that in and incorporated that into our guidance for the year in terms of the free cash flow margins that we expect to generate. So, it's just normal course for us. It's nothing different than what we had done previously with respect to the data centers.

Operator

Ladies and gentlemen, that concludes today's call. Thank you all for joining. You may now disconnect.

