

04-Nov-2024

# DigitalOcean Holdings, Inc. (DOCN)

Q3 2024 Earnings Call

## CORPORATE PARTICIPANTS

**Melanie Strate**

*Head-Investor Relations, DigitalOcean Holdings, Inc.*

**Padmanabhan T. Srinivasan**

*Chief Executive Officer & Director, DigitalOcean Holdings, Inc.*

**W. Matthew Steinfort**

*Chief Financial Officer, DigitalOcean Holdings, Inc.*

---

## OTHER PARTICIPANTS

**Raimo Lenschow**

*Analyst, Barclays Capital, Inc.*

**Mike Cikos**

*Analyst, Needham & Co. LLC*

**James E. Fish**

*Analyst, Piper Sandler & Co*

**Gabriela Borges**

*Analyst, Goldman Sachs & Co. LLC*

**Jeff Hickey**

*Analyst, UBS Securities LLC*

**Joshua Baer**

*Analyst, Morgan Stanley*

**Pinjalim Bora**

*Analyst, JPMorgan Securities LLC*

---

## MANAGEMENT DISCUSSION SECTION

**Operator:** Thank you for standing by and welcome to the DigitalOcean Third Quarter 2024 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. [Operator Instructions] Thank you.

I'd now like to turn the call over to Melanie Strate, Head of Investor Relations. You may begin.

---

**Melanie Strate**

*Head-Investor Relations, DigitalOcean Holdings, Inc.*

Thank you and good morning. Thank you all for joining us today to review DigitalOcean's third Quarter 2024 financial results. Joining me on the call today are Paddy Srinivasan, our Chief Executive Officer; and Matt Steinfort, our Chief Financial Officer. After our prepared remarks, we will open the call up to a question-and-answer session.

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. I refer specifically to the discussion of our expectations and beliefs regarding our financial outlook for the fourth quarter and full-year 2024, as well as our business goals and outlook. Our actual results may differ materially from those projected in these forward-looking statements. I direct your attention to the risk factors contained in our filings with the Securities and Exchange Commission and 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 measured on this conference call, and reconciliations to the most directly comparable GAAP financial measures are also available in today's 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

*Chief Executive Officer & Director, DigitalOcean Holdings, Inc.*

Thank you, Melanie. Good morning, everyone, and thank you for joining us today as we review our third quarter 2024 results. DigitalOcean had a successful third quarter, continuing to deliver progress on our key metrics and executing on the initiatives we laid out earlier in the year, further establishing ourselves as the simplest scalable cloud. In my remarks today, I will briefly highlight our third quarter results, share tangible examples of how our increased pace of innovation is benefiting our customers, discuss the continued momentum we are seeing with our AI/ML platform, and give an update on our strategic partnerships and engagement with the developer ecosystem.

First, I would like to briefly recap our third quarter 2024 financial results. Revenue growth remained steady in the third quarter at 12% year over year, with solid performance in core Cloud and continued growth in AI despite lapping difficult comps from our managed hosting price increase in April 2023 and from the Paperspace acquisition in July 2023. We continue to see momentum in demand for our AI/ML products, where Q3 ARR again grew close to 200% year-over-year.

In addition, we saw revenue growth, contributions from new customers and steady growth from our core business as we continue to enhance our customer success and go-to market motions.

Having delivered strong results through the first three quarters, we are increasing the lower end of our full-year revenue guide by \$5 million and the top end by \$2 million. We continue to focus the majority of our product innovation and go-to market investments on our builders and scalers who drive 88% of our total revenue and are growing 15% year-over-year ahead of our overall 12% revenue growth.

We also delivered strong adjusted EBITDA margins at 44% and have maintained our full-year free cash flow margin guidance as we continue to manage costs effectively while still investing to accelerate product innovation in Cloud and AI. Matt will walk you through more details on our financial results and guidance later in this call.

Let me start by giving you an update on our core cloud computing platform. In Q3, we continued our increased product velocity, specifically focused on the needs of our largest and fastest-growing customer cohort, the 17,000-plus scalers that drive 58% of our total revenue, and that grew 19% year-over-year in the quarter. In Q3, we released 42 new product features in total, which is almost double what we delivered in the previous quarter. We are accelerating features that will benefit our existing and potential scalers that are on other hyperscaler clouds today. Let me now provide a few highlights from these efforts that are specifically focused on the needs of these larger workloads.

We announced the early availability of Virtual Private Cloud peering, or VPC peering for short, that gives customers the ability to connect two different VPCs on the DigitalOcean platform within a data center or between different data centers. Through VPC peering, customers can create strong data isolation and privacy via direct and secure networking between resources that doesn't expose traffic to the public Internet.

Our Global Load Balancer or GLB is now generally available for all of our customers. GLB offers global traffic distribution based on geographical proximity of the end user, enabling lower-latency services; dynamic multiregional traffic failover, enabling more service availability for our customers' applications; data center prioritization; edge caching; and automatic scaling of the load balancers. We are thrilled to be able to roll it out to all of our customers, particularly to our scaler customers with existing multinational deployments that will benefit directly from this new product.

During the third quarter, we progressed daily backups from early availability to general availability, giving our customers the additional flexibility to manage backups at a daily and weekly cadence. This enables increased protection for our customers' workloads. As with daily backups, we automatically retain the seven most recent backup copies. This was an explicit need given the large volume and growth of data we are seeing on our platform with our Spaces Object Storage footprint growing 50% year-over-year.

We're also launching larger droplet configurations, including 48 vCPU memory and storage-optimized droplets, 60 vCPU CPU-optimized and general-purpose droplets, and larger 7-terabyte and 10-terabyte disk density variant droplets. These large droplet configurations are particularly relevant to our scaler customers who can quickly scale up their workloads that require more CPU, memory, or storage versus horizontally scaling out with multiple nodes.

In September, we announced Kubernetes Log Forwarding, which also enables Kubernetes customers centralized log management, simplifying the monitoring and troubleshooting of their applications in the DigitalOcean platform. This was built with simplicity in mind. With just a few clicks from the Kubernetes settings panel, customers can easily forward cluster event logs from Kubernetes directly to the DigitalOcean Managed OpenSearch for further analysis.

We also enhanced application security for our cloud-based Managed Hosting product by introducing a new malware protection solution and saw 3,650 net activations within the first week. To-date, we have seen near-zero false positives or false negative rates from our malware detection. This malware protection capability is now one of the fastest-growing revenue-generating product modules we have seen on our Managed Hosting platform.

All these innovations are not only helping us meet the needs of our large customers, but also helping us move customers with these larger workloads from purely usage-based to committed contracts. For example, an existing cybersecurity customer of ours, Siebel, a leader in threat intelligence, signed a multiyear seven-figure commitment in this quarter. The decision to continue leveraging DigitalOcean and sign a multiyear deal was driven by the release of our new large premium CPU-optimized droplet that helps customers run computationally heavy workloads. Siebel is a petabyte scale company and after several weeks of diligence, they chose DigitalOcean for this new workload due to our scalability, coverage, and cost efficiency.

Another great example is Traject Data who signed a multiyear commitment for a broad portfolio of DigitalOcean services, including over 500 droplets, Managed MongoDB, Spaces, Backups, and Volumes. Traject Data requires robust, scalable, and reliable infrastructure to power their real-time clean and bulk process data insights serving domains including marketing, retail, and analytics. They use the DigitalOcean platform to host their APIs and manage vast amounts of search engine results page and e-commerce data to deliver critical insights to their customers.

This product innovation and enhanced customer engagement is also helping customers migrate workloads to DigitalOcean from the hyperscalers. One specific example is Picap, a leading ride-sharing and logistics company

based in Latin America, operating in Mexico, Brazil, Peru and Colombia, and they moved all of their workloads from various clouds to DigitalOcean in the third quarter. They migrated to DO due to the simplicity of our products, transparent and simple pricing model, and strong support from our customer-facing teams.

Another example is [ph] NoBid (00:10:50), a customer specializing in optimizing ad revenue for online publishers through real-time bidding technology. Upon technical validation of the DO platform scale, they moved most of their large-scale production applications from a hyperscaler to the DO platform, reinforcing our opportunity to increase our share of wallet with our scaler customers.

Next, let me provide some updates on the AI/ML side. Our AI strategy reflects our belief that the AI market will evolve in a similar fashion to other major technology transformations, with initial progress and monetization at the infrastructure layer, which will eventually be eclipsed by the opportunities and value creation at the platform and application layers. Like others in the market today, we are actively participating in the infrastructure layer, but we are also innovating rapidly in the platform and application pillars to make it easy for our customers to use GenAI at scale without requiring deep AI/ML expertise. This is where we see our differentiation, as our customers seek to consume AI through platforms and agents rather than building everything themselves using raw GPU infrastructure.

At the infrastructure layer, we made GPU Droplets accelerated by NVIDIA H100 Tensor Core GPUs generally available to all of our customers. Now, all DigitalOcean customers can leverage on-demand and fractional access to GPUs, which is a critical step in achieving our overarching mission of democratizing AI for all customers.

In Q3, we also announced the early availability of NVIDIA H100 Tensor Core GPU worker nodes on the DigitalOcean Kubernetes platform, or DOKS for short, providing customers with a managed experience with GPU nodes ready with NVIDIA drivers, NVLink Fabric Manager, and NVIDIA Container Toolkit. Customers can take advantage of the NVIDIA GPU Operator and NVIDIA Mellanox network operator to install a comprehensive suite of tools required for production deployment. Both GPU Droplets and the H100 GPU nodes on DOKS are examples of how we are innovating even in the infrastructure layer, making it simpler for customers.

Let me give you an example. [indiscernible] (00:13:22) Exchange is a paytech company that specializes in providing enterprise blockchain-based solution for bank payments, and they are leveraging DigitalOcean's H100 infrastructure to accelerate the processing of high-volume financial transactions by providing advanced computational power. They use machine learning models to detect fraud in real time, assess risk, and ensure that payments are processed securely and quickly. The GPU infrastructure allows them to process more transactions while maintaining low latency and improving the overall user experience for both banks and end customers.

Next, at the platform layer, in this quarter we launched the early availability of our new GenAI platform to select customers so that we can iterate with them and shape the product and make it easy for them to build GenAI applications that deliver real business value. Users of this product will be able to combine their data with the power of foundational model to create personalized AI agents to integrate with their applications in just a few minutes. Customers can leverage our platform to create AI applications with foundational models and agent routing, knowledge bases, and retrieval-augmented generation, or RAG. This is a key step towards our software-centric AI strategy which is aimed at enabling customers to derive business value from AI in a friction-free manner.

An example of a customer that is already leveraging our GenAI platform is autonoma.cloud, a plant digitization company that offers a platform for manufacturing plants and machine manufacturers. autonoma.cloud creates and manages large volumes of documentation and data for each of their customers' plants and individual machines,

and we're looking to create AI agents that understood their user-specific contexts and retrieve answers and machine-specific data to their queries. With DO's new GenAI platform, they quickly built an interactive experience with their custom data that reduces the cognitive overhead for its users.

It is very important to note that these companies are not just doing internal proof of concepts or R&D projects, but are now starting to leverage our AI/ML products to build AI into their own products to deliver real business value to their customers without requiring deep expertise in AI, machine learning, data science or data engineering.

Finally, let me talk about the third pillar of our AI strategy, the application or agentic layer. As I just talked about, our customers are using our GenAI platform to create their own AI-driven agents. In addition to that, we are also innovating on this front by further simplifying cloud computing using AI and automating workflows that were previously done by humans.

One of the frequent pain points for our customers is debugging their cloud applications when something goes wrong, because, one, it is a very complex set of technical tasks; and number two, they typically don't have specialized site reliability engineers, or SREs, available in their staff to perform these complex tasks.

So, we set out to mitigate this pain point for our customers using GenAI by building a new AI agent to perform some of these tasks that are typically done by human SREs. We are using this AI SRE agent, both internally on our systems and externally by integrating it with our cloud-based product. Let me explain.

Internally, we are using the AI SRE agent to help our human SREs troubleshoot ongoing technical incidents in the DO Cloud platform. Based on our initial internal data, the AI SRE agent is reducing the time it takes to identify root causes by almost 35% by leveraging AI to quickly process an enormous amount of log data from disparate systems to pinpoint root causes and make next-step decisions, including recommendations to fixes or underlying problems.

Externally, we integrated this AI SRE agent into our cloud-based product, which hosts hundreds of thousands of mission-critical websites. Today, when issues happen on customers to servers and applications, they have to work with support engineers to debug the root cause then apply a fix. This is true not just for the DigitalOcean platform but across all managed hosting platforms. This can be a time-consuming job during which their business and even websites can be affected if not offline.

Our new AI SRE agent jumps into action upon detection of any performance degradation due to common issues like aggressive bot crawlers, denial of service attacks and so forth to investigate and gather insights and provide recommendations real-time on how to fix these issues, thereby reducing the time to resolution significantly. Our testing results are very encouraging, and we have just started working with a few customers in early availability mode.

Rounding out our AI strategy, we opened up a new front door by launching a strategic partnership with Hugging Face in Q3. Hugging Face is the leading open source and open science platform that helps users build, deploy, and train machine learning models. As a result of this partnership, DigitalOcean now offers model inferencing through one-click deployable models on GPU Droplets, allowing users to quickly and easily deploy the most popular third-party models with the simplicity of GPU Droplets and optimal performance accelerated by NVIDIA H100 Tensor Core GPUs.

This offering simplifies the deployment complexity of the most popular open-source AI/ML models, as DigitalOcean has natively integrated and optimized these models for GPU Droplets, enabling fast deployment and

superior performance. The Hugging Face partnership will make it easier for the more than 1.2 million Hugging Face users to discover and use the DigitalOcean platform.

In Q3, we also announced a new partnership with Netlify, a leading Web development platform to enable customers seamlessly connect their Netlify applications to DigitalOcean's Managed MongoDB, offering developers all the right tools to build and scale their applications without the complexities of managing infrastructure.

These announcements, in addition to the various other partnerships we already have in flight, highlight our efforts to augment our durable product-led growth motion with additional channels, including new front doors through partnerships with leading players in our ecosystem that will also help shape and improve our product offerings.

I'm also excited to highlight the material progress we are making with our renewed engagement with the developer community. In October, we hosted the 11th edition of Hacktoberfest, which has now evolved from being an internal hackathon event at DigitalOcean to one of the largest and premier opensource community events. This year, over 65,000 developers from 172 countries participated in more than 115 community run events and contributed to 15,000 open source projects.

Beyond Hacktoberfest, we also hosted more than 10 DigitalOcean meetups for developers and AI/ML community and participated in a number of industry conferences. This broad-based community engagement effort reinforces DigitalOcean's ongoing community to our developer ecosystem.

In closing, I am encouraged by the progress on product innovation and customer engagement, particularly as it is helping our builder and scaler customers continue to grow on our platform as their businesses expand. We are also making great strides towards our software-centric AI vision by rapidly shipping products in each of the three layers: infrastructure, platform, and applications. We're starting to see the green shoots from these investments in the form of customer wins, including cloud migrations from the hyperscalers, multiyear commitment contracts, and real-world deployment of AI using the DO AI platform. We will continue to focus on our largest and fastest-growing customer cohorts as we seek to accelerate growth in the quarters to come.

Before I turn the call over to Matt, I'm very excited to share that we will be hosting an Investor Day in New York City, and we are currently targeting late March or early calendar Q2 2025, in which we will share more on our long-term strategy including more detail on our progress and metrics, as well as a view of our long-term financial outlook.

I will now hand the call over to Matt Steinfort, our CFO, who will now provide some additional details on our financial results and our outlook for Q4 2024. Thank you.

---

## W. Matthew Steinfort

*Chief Financial Officer, DigitalOcean Holdings, Inc.*

Thanks, Paddy. Good morning, everyone, and thanks for joining us today. As Paddy just covered, we had a very successful Q3, both executing on key initiatives and delivering solid financial performance. In Q3, we continued to see increased momentum from our AI/ML platform and steady growth across our core business while consistently delivering attractive adjusted EBITDA and adjusted free cash flow margins.

Revenue in the third quarter was \$198.5 million, up 12% year over year. Annual run rate revenue, or ARR, in the third quarter was \$798.3 million, also up 12% year-over-year. We added \$17 million of ARR in the quarter. Most notably, builders and scalers, which are our largest customers, together grew 15% year-over-year.

Contributing to our overall growth was healthy incremental revenue from new customers and increased momentum from our AI/ML platform which saw significant growth, again growing close to 200% year-over-year on an ARR basis. Overall growth was partially muted by our Managed Hosting platform as we are lapping difficult comps related to the April 2023 Managed Hosting price increase and a temporal surge of Managed Hosting revenue in Asia in late 2023.

Our Q3 net dollar retention rate was steady at 97%. As with prior quarters, we're continuing to see consistent but below historical net expansion levels while our churn levels have remained low for well over a year. We will continue efforts to improve growth and NDR, including executing on our product road map and working to layer on additional go-to market motions to complement our durable product-led growth engine.

Turning to the P&L, gross margin for the quarter was 60%, which was 100 basis points lower than the prior quarter and consistent with the prior year. We are able to maintain healthy gross margins while continuing our investment in AI infrastructure due to the success of our ongoing cost optimization efforts.

Adjusted EBITDA was \$87 million, an increase of 14% year-over-year. Adjusted EBITDA margin was 44% in the quarter, approximately 200 basis points higher than the prior quarter. This increase quarter-over-quarter was primarily driven again by our ongoing operating cost discipline.

Diluted net income per share was \$0.33, a 65% increase year-over-year; and non-GAAP diluted net income per share was \$0.52, an 18% increase year-over-year. This increase is directly a result of our ability to increase our per-share profitability levels by continuing to drive operating leverage while mitigating dilution through share buybacks.

Finally, Q3 adjusted free cash flow was \$26 million or 13% of revenue. This is lower than the prior quarter by approximately 600 basis points due to timing of capital expense payments as we continue to make investments capitalized on the AI opportunity to fuel future growth. As a reminder, quarterly free cash flow margin will vary given the timing of capital spend and other working capital impacts. The lower free cash flow in Q3 does not change our expected full-year free cash flow margin.

Turning to our customer metrics. The number of builders and scalers on our platform, those that spend more than \$50 per month, was approximately 163,000, representing an increase of 6% year-over-year. The revenue growth associated with builders and scalers was 15% year-over-year, ahead of our overall revenue growth rate of 12%. The number of builders and scalers on our platform which, together, represent 88% of our total revenue increased by 2,260 quarter-over-quarter. The continued growth of our larger spending cohorts is a direct result of our focused product development, much of which is driven by direct customer feedback and the customer success and go-to market investment that are concentrated on these builders and scalers.

Our overall revenue mix continued to shift more towards our higher spend and higher growth customers, and we saw total ARPU increased 11% year-over-year to \$102.51.

Our balance sheet remains very strong as we ended the quarter with \$440 million of cash and cash equivalents. We also continue to execute against our share repurchase program with \$11 million of repurchases in the quarter, bringing total share repurchases to \$29.9 million through the first three quarters of the year.



With our healthy cash position and ongoing free cash flow generation, we are well-positioned to continue to balance investment in organic growth with share repurchases while moving towards our 2.5 to 3 times net leverage target and maintaining appropriate flexibility to address our 2026 convert at the appropriate time.

Moving on to guidance. Based on our performance year-to-date, we are increasing the bottom end of our full-year 2024 revenue guide by \$5 million and the top end by \$2 million, projecting revenue to be in the range of \$775 million to \$777 million, a \$3.5 million increase in the midpoint of our guidance range which will represent year-over-year growth of approximately 12%. This full-year guide implies Q4 revenue to be in the range of \$199 million to \$201 million, representing approximately 11% year-over-year growth at the midpoint of our guidance range. While we are not yet going to provide 2025 revenue guidance, we expect to enter 2025 with baseline growth in the low- to mid-teens.

As demonstrated throughout 2024, we remain committed to driving continued operating leverage in our core DigitalOcean platform. Given our solid performance throughout the first three quarters of the year, we are raising our adjusted EBITDA margin guidance for the full year to be in the range of 40% to 41%. This full-year adjusted EBITDA guide implies Q4 adjusted EBITDA margins to be in the range of 34% to 38%.

For the full year, we expect non-GAAP diluted earnings per share to be \$1.70 to \$1.75. This implies Q4 non-GAAP diluted earnings per share to be \$0.27 to \$0.32 based on approximately \$103 million to \$104 million in weighted average fully diluted shares outstanding.

Turning to adjusted free cash flow, we expect adjusted free cash flow margins for the full year to be in the range of 15% to 17%, consistent with what we guided in the prior quarter. While free cash flow margin will continue to vary quarter to quarter, we anticipate remaining at a similar 15% to 17% range on a rolling average quarterly basis in 2025 as we continue to accelerate the pace of product innovation and make disciplined investment to expand our emerging AI capabilities.

That concludes our prepared remarks, and we will now open up the call for Q&A.

## QUESTION AND ANSWER SECTION

**Operator:** Thank you. We will now begin the question-and-answer session. [Operator Instructions] Your first question comes from the line of Raimo Lenschow from Barclays. Your line is open.

**Raimo Lenschow**

*Analyst, Barclays Capital, Inc.*

Q

Perfect. Thank you. Paddy, there was a lot of product innovation that you kind of discussed. Can you talk a little bit about how we have to think about those new innovations around products and how that feeds into the installed base in terms of like what's the uptake there, what's the timing there because the financial number Matt mentioned also look at it as the NRR is 97% stable, but there seems to be a little bit of a disconnect. Can you maybe talk to kind of timing there, etcetera.? And I have one follow-up for Matt.

**Padmanabhan T. Srinivasan**

*Chief Executive Officer & Director, DigitalOcean Holdings, Inc.*

A

Well, thank you, Raimo, very much for the question. Yes, we are we are seeing a lot of product innovation across the board, both in the core Cloud. That's why I spent so much time explaining all the things we are pumping out, especially aimed at our scalers and allowing them or enabling them to run larger workloads on DigitalOcean.

As you know, from a timing and sequence point of view, there's no magical answer that we can provide, which translates product innovation to adoption and, hence, impact on our financial performance. But we see we have to do this to enable our customers move many of their larger workloads that they're currently running in other clouds and make it super compelling for them to run those workloads on the DO platform.

And as I did just a few minutes ago, we will get into a habit of explaining some very concrete examples of customers that are starting to do that. So, the examples I gave, we are now starting to sign customers into multiyear contracts with commitments on our platform. We are also starting to see a steady dose of migrations coming from other clouds, especially the hyperscalers.

So, we have to ensure that we have patience in terms of building these capabilities. We are starting to see the green shoots in terms of customer adoption and translation of that into leading indicators. And I have no question, if we keep doing it for a handful of quarters, we are going to start seeing the translation into other lagging indicators, including some of the ones that you just mentioned, Raimo. And...

**Raimo Lenschow**

*Analyst, Barclays Capital, Inc.*

Q

Okay. That's – oh, yeah, go ahead. Sorry.

**Padmanabhan T. Srinivasan**

*Chief Executive Officer & Director, DigitalOcean Holdings, Inc.*

A

No. So in terms of the NDR, I think Matt alluded to this fact, that we have – what we are seeing from a core – the NDR of the core business is trending a little bit ahead of what we are reporting on a blended basis. So, it gives us enough reasons to believe that what we are doing is starting to be appreciated by our customers. And as you know, this takes time for the adoption to happen. I have to keep reminding ourselves that we have 638,000-plus customers, so it takes time for the propagation to happen across the board with our customer base.

**Raimo Lenschow**

*Analyst, Barclays Capital, Inc.*

Q

Okay. Perfect. Thank you. And then, Matt, one for you. Like if you look on the EBITDA, that's kind of the one where you're kind of, at the moment, outperforming quite a bit. Can you talk a little bit about that? How do you keep – like how sustainable is the progression there, especially as you think about like more services coming on stream, you probably want to support them more, and then obviously more AI services coming as well. Thank you.

**W. Matthew Steinfort**

*Chief Financial Officer, DigitalOcean Holdings, Inc.*

A

Thanks, Raimo. Yeah, I'd say from a cost standpoint, Q3 was definitely a good quarter from an EBITDA margin perspective. As we brought on our new executives, we had talked about, implicit in our guide for the full year, we were making sure that we had enough room to invest to enable them to really improve the pace of innovation and layer on additional go-to-market motions. But at the same time, we were evaluating, okay, what costs do we have now that we just aren't earning a return on and can we clean those out before the team gets going with the new expenses.

So, we, I think, did a really good job of optimizing for that, and we also made some decisions to make sure that we were appropriately pacing the increases to see if we're getting a return on the investments as we did it. So, I'd say, it was just disciplined kind of cost management in Q3. And as you saw from the guide in Q4, we are expecting to ramp our expenses heading into next year. I don't think it's going to be a meaningful kind of a change in the overall expense level. We feel pretty good about kind of the trailing margin profile that we have and being able to continue that into next year.

**Raimo Lenschow**

*Analyst, Barclays Capital, Inc.*

Q

Okay. Perfect. Thank you.

**Operator:** Your next question comes from the line of Mike Cikos from Needham. Your line is open.

**Mike Cikos**

*Analyst, Needham & Co. LLC*

Q

Hey, thanks for taking the question, guys. I think the first would go to Matt just coming off of Raimo's question there. But if I look at the EBITDA guide that we have today, the 34% to 38% margin guide in Q4 is the widest range that I think we've had in recent memory. And just wanted to get a little bit more granular there as far as, I guess, what needs to go wrong or right or what you guys are weighing for that gets you at the 34% margin versus the 38% margin in the December quarter?

**W. Matthew Steinfort**

*Chief Financial Officer, DigitalOcean Holdings, Inc.*

A

Yeah. It's a good question, Mike. I think a part of it is as we've been ramping particularly on the R&D spend, we're evaluating kind of surge resources using contractors to accelerate a handful of things on the product road map, and the timing of that, which, again, I view that as a relatively lumpy potential investment and the timing of being able to get that spun up and fully staffed and moving, whether that hits in Q1 or whether it hits in Q4, I think that's really what's causing the range.

Again, I think on a go-forward basis, we don't anticipate a material change in the overall kind of R&D as a percent of revenue, but we are advancing the expense. So, in any one quarter, it may be a little bit lumpier. But again, over a longer period of time, we don't – we think we'll grow into that, and some of that is surge resource.

**Mike Cikos***Analyst, Needham & Co. LLC*

Q

Terrific. And just another follow-up, I know that you guys aren't providing explicit guidance for calendar 2025 here. I do appreciate the qualitative commentary. Just wanted to see what gives you the confidence to kind of put that bogey out there for the baseline growth and how should we be thinking about what it takes for DigitalOcean to be entering the year with that kind of baseline growth that you had commented on.

**W. Matthew Steinfort***Chief Financial Officer, DigitalOcean Holdings, Inc.*

A

Yeah. I think it's very similar to what we described at the beginning of this year, right? What can you count on? Well, you count on the growth from the self-serve funnel, and we're a little bit better – doing a little bit better on that year-to-date than we had outlined at the beginning of the year. We've got the Managed Hosting business, which is kind of returning to growth after lapping some difficult comps. We've got AI/ML that we had said would contribute around 3 points of growth. It's a little bit ahead of that for the year. And then NDR, while it's frustrating that we've had to print a bunch of 97s in a row, as Paddy said, the core DO NDR is actually ahead of that. It's a little bit of a headwind for Managed Hosting that's going to be in place through the next – call it, through the first quarter of next year.

And so, if you take all those together, we've moved up a couple of points from the baseline growth that we had described coming into the year, and none of it is on the back of kind of macro improvements. It's all just steady kind of improvements and continuing to deploy products that our customers need. And so, as we look at that trajectory, we feel comfortable kind of at the pace of growth that we're at right now and, hopefully, continuing to improve NDR every month going into next year and beyond to eventually get it to be above 100%. So, I'd say we're just making sure that folks understand that we do feel pretty comfortable with the baseline growth that we're delivering.

**Mike Cikos***Analyst, Needham & Co. LLC*

Q

Great answer. Thank you very much, guys.

**Operator:** Your next question comes from the line of James Fish from Piper Sandler. Your line is open.

**James E. Fish***Analyst, Piper Sandler & Co*

Q

Hey, guys. Paddy, for you, you guys are seeing adoption – are you seeing adoption of the GPU Droplet with more of the builders and scalers or more net new customers and how should we think about the mix between on-demand versus multiyear contracts and what you guys are seeing around supply availability of the GPUs?

**Padmanabhan T. Srinivasan***Chief Executive Officer & Director, DigitalOcean Holdings, Inc.*

A

Yeah. Great. Thank you, Jim, for the question. So in terms of the adoption, we are seeing adoption across the board. A lot of new customers, which we absolutely love that are kicking the tires and also, as I explained on the

call, building real-world applications on our GPU infrastructure, both Droplets as well as more hardened bare-metal type of services, I would say from between on-demand versus contract, we see more contracts when the customer is deploying live workloads, whether it is training or inferencing. And sometimes these contracts are fairly short term but some are longer term. And on-demand is typically for experimentation, which is what we would have guessed when we started this journey, but that's where things are.

And from an on-demand point of view, we're also seeing a very nice uptake and interest in our GenAI platform. So, companies that don't have the deep bench in terms of AI/ML skill sets have a very easy time just using our GenAI platform, standing something up very quickly many times in just a matter of few minutes, just to see if they can prove to themselves that there's value in integrating GenAI into their platform. So, that's what we are seeing broadly from a adoption point of view.

---

**W. Matthew Steinfort**

*Chief Financial Officer, DigitalOcean Holdings, Inc.*

**A**

Then from a supply chain standpoint – yeah, on the supply chain, we don't see the same kind of headwinds that we had seen coming into the year. We've got orders out for the next generations of the technology. The H200 is coming. We're keeping an eye on Blackwell the timing of that. And it's certainly not so tight that you can get it in a week or two from ordering, but it's our data supply chain is open enough that we've been able to get the equipment in the time frames that we need it. And again, with our buildout of the Atlanta data center coming on kind of at the beginning of next year, we're in good shape from a logistics and a scheduling standpoint.

---

**James E. Fish**

*Analyst, Piper Sandler & Co*

**Q**

Got it. And then, Matt, for you, circling back on the 97% net expansion rate, the AI side of things turned organic this quarter. By my math, it's probably adding about 1 to 2 points to NRR. It looks like net new ARR for AI was up around \$10 million. So, what's going on with that core business? Specifically, you were starting to mention around Cloudways. Obviously, the price increase lapping but why is that business kind of weaker than what you guys are anticipating and how should we think about the mix of Cloudways hosting DigitalOcean versus other cloud platforms?

---

**W. Matthew Steinfort**

*Chief Financial Officer, DigitalOcean Holdings, Inc.*

**A**

That's a great question and good clarification. The AI products are not in NDR. So, just to make sure that everybody understands that. The revenue from the AI products are not in net dollar retention. And that's – it's clear in the definitions that we have – that lot of the AI revenue, if you think about it, is project-based, right? So someone's training something, somebody is coming in and experimenting, it's not yet at the point where people are coming in and running large-scale inference workloads where you could say, oh, well, that should – the revenue that you get from that inference workload should be bigger next year than it was this year because they're spending – they have a lot more customers. If someone comes in and train the model for a month or two and then turns it off and then does – goes and focuses on inferencing, the revenue is going to be lumpy. And so, at this point – and we could re-evaluate this going forward. But at this point, AI is not reflected in NDR, so contributes nothing to the improvement in the NDR that we've seen.

What we have seen is steady improvement in the core Cloud business, which we said is tracking above the reported NDR. Cloudways, which has historically been literally until we lapped the price increase. It was always a positive contribution to NDR. It's been a headwind to NDR since April and will continue to be probably until next April because of the just [ph] vagaries (00:44:17) of a lagging metric like NDR. But we expect both the

Cloudways, the Managed Hosting business and the core DO business, we expect to be able to get those back above 100% and we're certainly working aggressively to accomplish that, and we can't tell you exactly when that's going to happen, but we're very encouraged by the green shoots that we're seeing in both businesses on that improvement in NDR.

---

**James E. Fish**

*Analyst, Piper Sandler & Co*

Q

Very helpful. Thanks, Matt.

---

**Operator:** Your next question comes from the line of Gabriela Borges from Goldman Sachs. Your line is open.

---

**Gabriela Borges**

*Analyst, Goldman Sachs & Co. LLC*

Q

Hey. Good morning. Thanks for taking the question. Matt, I wanted to follow up on some of your comments for 2025 and more specifically on how we should think about the seasonality of the business given some of the moving pieces we've had this year versus last year, so any comments on seasonality.

And I'm noticing that the size of the beat this quarter was about 1% versus the 2% last quarter. Any nuances we should be aware of there in terms of why the size of the beat was smaller this quarter? Thank you.

---

**W. Matthew Steinfort**

*Chief Financial Officer, DigitalOcean Holdings, Inc.*

A

I don't think there's any seasonality in the business that would reflect that. I think that, again, we've been very focused on the full year and providing guidance that's appropriate and reflective of that. And that causes kind of bigger swings in the quarterly kind of beats rates where we're more focused, I think, from an annual standpoint.

But I think that as we look at the business going into next year, again, going back to my earlier comments, we're very encouraged by the steady growth that we're seeing in improving growth versus what we thought with the self-serve funnel and feel comfortable about that. The Managed Hosting business is coming back from, again, some difficult comps. The AI business is slightly ahead of where we had expected. And then kind of the last thing to move for us, which would give us the confidence to increase our outlook on the revenue is that NDR just needs to come up and expand steadily. Steadily but stubbornly moving up. And so, I don't think there's anything seasonal that would suggest we would beat more or less on an individual quarter.

---

**Gabriela Borges**

*Analyst, Goldman Sachs & Co. LLC*

Q

Got it. Okay. And then the follow-up is for Paddy. So given the pace of change that we're seeing in the GPU-as-a-service market, maybe you could walk us through what are one or two of the areas where you feel like you've learned the most over the last three months as it relates to your AI service strategy and particularly around your LLM-as-a-service offerings, the platform offerings, how you think you can differentiate versus something like a SageMaker or a Bedrock over the long term. Thank you.

---

**Padmanabhan T. Srinivasan**

*Chief Executive Officer & Director, DigitalOcean Holdings, Inc.*

A

Thank you, Gabriela. Great question. So in terms of what we have learned over the last 90 days, we have learned a lot. As you can see, we have also shipped a lot. So in preparation of that, I think we have learned quite a bit on

all three layers of our platform. I would say for me, personally, the biggest learning has been that our customers, which are typically companies that don't have a tremendous bench of deep machine learning, data scientists, or data engineering skill sets, they look at the AI platform almost in an inverted fashion.

What I mean by that is everyone, us included, the market, everyone looks at it from infrastructure first and then platform and then finally applications. Our customers actually look at it top down. They look at, okay, what applications can I or agents can I leverage today from GenAI that makes my app more productive or my customers save money or deliver some innovation that was not possible so far. So, it's almost a realization that we need to innovate more rapidly on the platform application layer. It's why we accelerated some of our GenAI platform capabilities, and we already have seen a customer push that into production which is amazing.

And the part two of your question is what makes our GenAI platform stand out against something like a SageMaker or Bedrock. As you know, we have very deep expertise in both SageMaker and Bedrock at DigitalOcean today. And the biggest difference is some of the technologies you mentioned are phenomenal. They're very broad and very powerful if you have a broad set of skills available to take all of that and build something fantastic for a very complex use case.

For our customers, and the customer that I talked about during the prepared remarks, specifically tested a variety of different GenAI platforms and picked us primarily because of how easy it was for them to get started to inject their own custom data to build a rack pipeline, to create a knowledge base, and finally create a chatbot where they could project exactly how much it would cost them and develop a business model that would be friendly to their customers.

So, all of these things individually are fairly complex. But when you add these different steps to build a productionized application, it just balloons in its complexity. And we have tried to measure every click it takes to simplify the journey for our customers, and I think that's how we establish ourselves as a credible cloud provider, and that's what we are doing to establish ourselves and differentiate ourselves in GenAI.

And also, we should also not forget that there is a lot of differentiation we are pushing even in the agentic layer. As I explained, we just came out with our first agent. We are working with customers in early availability mode, so we will learn and innovate on that faster. But the combination of the platform and application gives us the ability to make things that are super scalable but, at the same time, an order of magnitude simpler to use compared to other alternate platforms that are available.

---

**Gabriela Borges**

*Analyst, Goldman Sachs & Co. LLC*



Thank you for the detail.

---

**Operator:** Your next question comes from the line of Jeff Hickey from UBS. Your line is open.

**Jeff Hickey**

*Analyst, UBS Securities LLC*



Hey, everyone. Thanks for taking the question. The first one I wanted to ask is it's very helpful just detailing that AI is not included in the NDR metric. But maybe with some of the existing AI customers that you've had for a few quarters, that maybe do have some workloads already in production, do you have any sense of like how they're expanding their spend over time, maybe even just on a quarterly basis or do you typically see those customers kind of launch a workload and then have that spend at sort of a stable level from there?

**W. Matthew Steinfort**

*Chief Financial Officer, DigitalOcean Holdings, Inc.*

A

We've seen good traction with a number of our early AI customers that have come in and experimented on the platform, and they may have started with kind of a small cluster. And then as they tested, they've expanded their use of the platform. So, the question is, when we land customers, do we see them grow or do we see a big rotation of customers in and out? We actually see fairly healthy expansion from the customers when they come in. But again, back to my earlier comment, it's an – okay, I'm training a model, I need eight nodes and now I'm going to do something, I need more. But it's not the same dynamic because they're still evaluating, they're still kind of going through the testing phases. But we've seen very good traction, growing customers, the initial customers that we've had on the AI platform.

**Padmanabhan T. Srinivasan**

*Chief Executive Officer & Director, DigitalOcean Holdings, Inc.*

A

And, Matt, one thing that I will add to that is it's interesting to note that our AI customers are also very similar to our core cloud customers, in the sense that most of them, if not all of them, are ISVs, or independent software vendors, or digital-native application providers. So, they are taking our – they're building solutions on our AI platforms, whether it is infrastructure or GenAI, to create software solutions for their customers. So, as they grow and expand, they will – they are expanding their footprint, to Matt's point, on our platform. So, that's a very interesting thing for us to notice versus a customer that is coming to build a solution just for their internal use.

**Jeff Hickey**

*Analyst, UBS Securities LLC*

Q

Got it. That's really helpful. And then one just quick follow-up. You've mentioned earlier about just supply, and that's gotten better relative to the beginning of the year with AI investments. Just curious, with the October 1 launch of H100 instances broadly available, are you supply-constrained at all right now as we're kind of in the fourth quarter or are you able to meet all the demand you currently have as well? Thanks.

**W. Matthew Steinfort**

*Chief Financial Officer, DigitalOcean Holdings, Inc.*

A

We've ordered enough – and we talked about this in the last earnings call that because we have the ability to see the demand and plan out the capacity, that we've been able to get enough capacity to meet the demand as we've gone, which is a very good sign, because we don't have those supply constraints. So, again, because we're not spending hundreds of millions of dollars on GPU, we can get the quantum that we need to meet our requirements. And when you have something like the GPU Droplet, which is more on-demand than a less committed contract, you have to see kind of what the utilization is, and then plan your purchases based on that capacity utilization, and we've been able to manage that effectively. So, it hasn't been a drag or a constraint to us.

**Jeff Hickey**

*Analyst, UBS Securities LLC*

Q

Got it. Very helpful. Thank you.

**Operator:** Your next question comes from the line of Josh Baer from Morgan Stanley. Your line is open.

**Joshua Baer**

*Analyst, Morgan Stanley*

Q



Great. Thanks for the question. One for Paddy. I guess, just thinking about the 42 new product features more than prior period and I think calling out some of them features that hyperscale customers are generally looking for, moving contracts to committed contracts, even migrating workloads from hyperscalers to DO, it's like in the past the story was more about DigitalOcean's simplicity of the platform and better support, lower pricing and maybe a little bit less about sort of getting into the competitive dynamic with hyperscalers. Just wondering, is it a right takeaway that there is a little bit of a shift in focus either up market or a little bit of expansion outside the simplest startups and SMBs just to be a little bit more competitive in the market. Is there a strategy shift there?

---

**Padmanabhan T. Srinivasan**

*Chief Executive Officer & Director, DigitalOcean Holdings, Inc.*

A

Yeah. Thank you, Josh. Great question as always. The shift is essentially following our customers' lead, honestly. So, as I made a point during the prepared statements to make sure that we are not abandoning or taking our eye off the importance of DigitalOcean and the developer community, we continue to nurture that. In fact, we are doing a lot more with the developer ecosystem this year compared to the recent past. But at the same time, we do recognize that we have 17,000-plus scalers who, on an average, spend more than \$25,000 with us. That's big, and that's 58% of our revenue. And if you add scalers, that's almost 88% of our revenue, which are growing much faster than our blended average growth rate. So, we have a unique opportunity to follow their leads and make sure that we are delivering capabilities that will enable them to run or expand their footprint on DigitalOcean.

We are increasingly in a multi-cloud world, even for smaller customers like the ones we target, and there is an opportunity for us to keep expanding our share of wallet with these companies. And the examples that I shared are just the starting point for what we believe our fair share slice of this enormous market. And if we keep doing what we are doing now, which is add compelling feature sets that enable our scaler customers to expand their footprint on us, I think there's a lot of value to be created for our customers on our platform.

---

**Joshua Baer**

*Analyst, Morgan Stanley*

Q

Very clear. If I could follow up with one for Matt. Just on the – like, some of the factors called out, the Managed Hosting, tough comp, pricing increases, the Asia influx of revenue, even some of the M&A, like get how that could be impacting some of the – like the net dollar retention rate or the year-over-year growth. If I'm just looking at quarter-over-quarter net new ARR out of \$32 million last quarter and \$17 million this quarter, anything to call out as far as that difference just on a quarter-over-quarter basis?

---

**W. Matthew Steinfort**

*Chief Financial Officer, DigitalOcean Holdings, Inc.*

A

Yeah. It's a good point, Josh. The big difference was the availability, like we brought on a ton of AI capacity in Q2, which we had pent-up demand for. So, we got a bump, a material bump in ARR last quarter. If you look at, we were around \$17 million or \$18 million in the quarter before and then we jumped to \$30 million. Now, back on to \$17 million. I'd say last quarter was more of an anomaly than this quarter. Clearly, we're looking to add more incremental ARR going forward, but most of that change was a result of a surge in AI capacity last quarter.

---

**Joshua Baer**

*Analyst, Morgan Stanley*

Q

Got it. Thank you.

---

**Operator:** Our next question comes from the line of Pinjalim Bora from JPMorgan. Your line is open.

**Pinjalim Bora***Analyst, JPMorgan Securities LLC*

Great. Thanks, guys. Just one question from me. The baseline growth outlook that you kind of shared entering 2025 seems pretty positive. It calls for an acceleration from negative growth rate in this year. So, I wanted to understand that a little bit more. Are you seeing some signals that projects will accelerate next year around the core business based on some of the customer conversations? Does that assume 100% NDR as you – as we go into Q1 and how should we think about AI [indiscernible] (01:00:05)?

**W. Matthew Steinfort***Chief Financial Officer, DigitalOcean Holdings, Inc.*

Pinjalim, I didn't hear the last part of that question. I heard the first part. But let me answer and then you can maybe come back with the AI question. We're seeing a lot of green shoots around kind of like, as you said, all the product traction that we're getting and some of our larger customers even being willing to commit to long-term contracts and commitment contracts, which isn't something that the company has done extensively in the past.

But then, core NDR is improving steadily. We're not assuming that it gets to 100% by Q1. That's not implicit in the kind of comments that we made regarding next year. I mean, we're going to work aggressively to get it to be 100%, but we can deliver the growth rates that we talked about because we're effectively delivering that now, right, at 12% growth with an NDR that's only 97%.

And we expect both the Managed Hosting NDR and the core Cloud NDR to improve as we head into next year, and we'll continue to get growth – very positive growth contributions from our AI capabilities. We said earlier this year that we thought we'd get 3% of the overall growth from AI, and we're a little bit ahead of that this year. So that's also positive and encouraging as we think about what the baseline growth is heading into next year.

**Pinjalim Bora***Analyst, JPMorgan Securities LLC*

Understood. One quick follow-up. The multiyear commitment is definitely interesting. Are you leaning in on any way to drive those commitments? Is that largely coming from customers or are you putting in processes to kind of enable those discussions? Thank you so much.

**Padmanabhan T. Srinivasan***Chief Executive Officer & Director, DigitalOcean Holdings, Inc.*

Yeah, I can take that. So, we are – at this point, Pinjalim, we're just letting it happen organically, so we don't have any pronounced establish go-to-market motion. We're not pushing it on our customers. We're just letting it organically happen. The most important thing for us is to learn, learn the patterns, learn what kind of technologies we need to build, learn the migration process itself and things like that. And going into next year, we, of course, will look into packaging it a little bit, productizing it, and also expand our third-party ecosystem that can help orchestrate some of these things.

So, there's a lot of work to be done to scale it. But right now, we are focused on nailing it and understanding exactly what it takes to be successful.

**Pinjalim Bora***Analyst, JPMorgan Securities LLC*

Thank you.

**Operator:** And that is all the time we have for questions. This concludes today's conference call. Thank you for your participation. You may now disconnect.

**Disclaimer**

The information herein is based on sources we believe to be reliable but is not guaranteed by us and does not purport to be a complete or error-free statement or summary of the available data. As such, we do not warrant, endorse or guarantee the completeness, accuracy, integrity, or timeliness of the information. You must evaluate, and bear all risks associated with, the use of any information provided hereunder, including any reliance on the accuracy, completeness, safety or usefulness of such information. This information is not intended to be used as the primary basis of investment decisions. It should not be construed as advice designed to meet the particular investment needs of any investor. This report is published solely for information purposes, and is not to be construed as financial or other advice or as an offer to sell or the solicitation of an offer to buy any security in any state where such an offer or solicitation would be illegal. Any information expressed herein on this date is subject to change without notice. Any opinions or assertions contained in this information do not represent the opinions or beliefs of FactSet CallStreet, LLC. FactSet CallStreet, LLC, or one or more of its employees, including the writer of this report, may have a position in any of the securities discussed herein.

THE INFORMATION PROVIDED TO YOU HEREUNDER IS PROVIDED "AS IS," AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, FactSet CallStreet, LLC AND ITS LICENSORS, BUSINESS ASSOCIATES AND SUPPLIERS DISCLAIM ALL WARRANTIES WITH RESPECT TO THE SAME, EXPRESS, IMPLIED AND STATUTORY, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, ACCURACY, COMPLETENESS, AND NON-INFRINGEMENT. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, NEITHER FACTSET CALLSTREET, LLC NOR ITS OFFICERS, MEMBERS, DIRECTORS, PARTNERS, AFFILIATES, BUSINESS ASSOCIATES, LICENSORS OR SUPPLIERS WILL BE LIABLE FOR ANY INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL OR PUNITIVE DAMAGES, INCLUDING WITHOUT LIMITATION DAMAGES FOR LOST PROFITS OR REVENUES, GOODWILL, WORK STOPPAGE, SECURITY BREACHES, VIRUSES, COMPUTER FAILURE OR MALFUNCTION, USE, DATA OR OTHER INTANGIBLE LOSSES OR COMMERCIAL DAMAGES, EVEN IF ANY OF SUCH PARTIES IS ADVISED OF THE POSSIBILITY OF SUCH LOSSES, ARISING UNDER OR IN CONNECTION WITH THE INFORMATION PROVIDED HEREIN OR ANY OTHER SUBJECT MATTER HEREOF.

The contents and appearance of this report are Copyrighted FactSet CallStreet, LLC 2024 CallStreet and FactSet CallStreet, LLC are trademarks and service marks of FactSet CallStreet, LLC. All other trademarks mentioned are trademarks of their respective companies. All rights reserved.