

08-Aug-2024

# DigitalOcean Holdings, Inc. (DOCN)

Q2 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

**Gabriela Borges**

*Analyst, Goldman Sachs & Co. LLC*

**Raimo Lenschow**

*Analyst, Barclays Capital, Inc.*

**Kingsley Crane**

*Analyst, Canaccord Genuity LLC*

**Pat Walravens**

*Analyst, Citizens JMP Securities LLC*

**Mike Cikos**

*Analyst, Needham & Co. LLC*

**Pinjalim Bora**

*Analyst, JPMorgan Securities LLC*

---

## MANAGEMENT DISCUSSION SECTION

**Operator:** Thank you for standing by, and welcome to the DigitalOcean Second 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 afternoon. Thank you all for joining us today to review DigitalOcean's second 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 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 third 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 discussed 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 for today's call is also available in the IR section of our website. And with that, I'll turn the call over to Paddy.

---

## Padmanabhan T. Srinivasan

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

Thank you, Melanie. Good afternoon, everyone, and thank you for joining us today as we review our second quarter results. DigitalOcean delivered a strong second quarter, building on the momentum from the first quarter and continuing to execute on all key metrics. In my remarks today, I will briefly highlight our second quarter results, provide an update on the leaders we hired recently, share tangible examples of our increasing product velocity, and discuss how we are capitalizing on our AI growth opportunity.

First, I would like to briefly recap our second quarter 2024 financial results. Revenue growth has continued to reaccelerate in the second quarter to 13% year-over-year, reflecting the growing signs of success we're seeing from both a product and go-to-market standpoint, and the continued acceleration of our AI and machine-learning products, where ARR has grown over 200% year-over-year from the Paperspace ARR we acquired last year.

In Q2, we also saw the largest step-up in incremental total company ARR in nearly two years, excluding the quarter in which we acquired the AI/ML business. We also delivered strong adjusted EBITDA margins at 42% and adjusted free cash flow margins at 19%, exemplifying our ability to demonstrate ongoing cost discipline and optimization while continuing to accelerate product innovation. Our second quarter financial results highlight the progress we are making and our ability to execute on the plans we laid out at the beginning of the year.

We're also encouraged by the signs of improvement within both our growth profile and our key fundamentals. Net dollar retention was flat versus the previous quarter at 97%, as expansion within our customer base continues to be lower than historical level, given we're still navigating a challenging macro environment, which is muting the positive impact of our increased product velocity and the stability we have seen in churn and contraction from our solid execution on the various customer success motions.

In addition to the increased momentum from our AI/ML products, we received healthy revenue contributions from both our managed hosting products and new customers. Matt will walk you through more details on our financial results and guidance later in the call.

In addition to our solid financial performance and accelerating product innovation traction, I'm also excited about the advancements we made in building out the team. We added three critical new leaders to our executive team over the past several weeks. First is Bratin Saha, our Chief Product and Technology Officer, who will lead product strategy, product engineering, infrastructure, and security. Most recently, Bratin built AWS's multibillion-dollar AI, machine-learning, and data platforms, which together represented one of its fastest-growing business segments. Previously, Bratin worked at NVIDIA and Intel, running many of their software infrastructure platforms.

We also announced Wade Wegner as our Chief Ecosystem and Growth Officer, which is a unique role that is highly appropriate for DigitalOcean as we are a very unique company. Our cost-efficient self-service customer acquisition model is one of the most efficient in the industry. As I have said many times, one of DigitalOcean's strengths and a key driver of our customer acquisition model is our passionate community of developers, many of whom have grown or are growing up learning to code on our platform. Wade and his organization will be responsible for supercharging our engagement with this community and for driving our very distinct product-led growth motion.

Finally, we recently announced Larry D'Angelo as our Chief Revenue Officer, who will bring his years of experience building and scaling high-velocity go-to-market teams to drive direct sales and partner sales to augment our product-led growth engine, and also to build scalable customer success and support functions to help our customers be successful and expand their footprint on our platform.

DigitalOcean's strong fundamentals and future potential drew these three world-class executives to come join us in our journey. Their arrivals have also created further hiring momentum, as having top talent such as these three new executives tend to attract additional world-class talent. We're already seeing this dynamic play out as they fill out their respective teams. I'm very confident that we now have the right executive team in place to fuel growth, increase product velocity, help our customers be successful, and to continue to execute on our mission of making cloud and AI simple and accessible for developers.

Now, let me give you an update on our products. As we continue to listen to our customers and incorporate their feedback, enabling them to grow and scale on our platform, we released 24 new product features throughout Q2, doubling our product velocity from the prior six months. We also revived Deploy, our virtual developer conference, which was held on July 9. I'm thrilled about the success of this event and look forward to continuing to engage with our developer community, as we intend to increase the frequency of our Deploy events and do them on a regular cadence going forward.

During our July event, we announced a number of material product announcements in response to customer feedback. First, we announced GPU droplets in early availability mode, and this launch democratizes on-demand access to NVIDIA H100 GPU instances for our customers, enabling them to leverage one, eight, or more GPUs at a time, providing flexible deployment options tailored to the various use cases and budgets. A lot more on this a little later.

During Deploy, we also announced our Global Load Balancer product, which we refer to as GLB, which is currently in public beta. This is engineered to bolster application resiliency, eliminate single points of failure, and significantly minimize end user latency and secure GLB traffic from denial-of-service attacks. It offers global traffic distribution based on geographical proximity of the end user, dynamic multiregional traffic failover, data center prioritization, edge caching, and automatic scaling of the GLBs. It is intuitive, predictably priced, and tailored to the essential needs of growing technology companies for enhancing their global resiliency.

We also recently announced that select DigitalOcean products can now be used to host electronic protected health information. This allows companies such as telehealth providers, healthcare software applications, and health tech organizations to build and scale sensitive workloads regulated under HIPAA on our developer cloud, leveraging select DigitalOcean-covered products.

During the quarter, we also launched Managed OpenSearch, a comprehensive solution designed for in-depth log analysis, simplifying troubleshooting, and optimizing application performance. With Managed OpenSearch, customers can now pinpoint and analyze log data with a lot of ease, customize log retention, enhance security of their applications, scale to fit capacity needs, and forward these logs from multiple sources.

During Q1, we announced that we offer Premium Memory-Optimized Droplets and Premium Storage-Optimized Droplets. And in Q2, we finished rolling this out to all of our data centers, and this was a huge milestone for us. We also announced improvements to our ad platform, including auto scaling, Dedicated Egress, and an expanded lineup with entry-level dedicated instances, higher data transfer allowances, and reduced bandwidth overage fees.

Dedicated Egress provides application developers with fixed IP addresses, enabling them to meet the security needs of their customers who run applications that require whitelisting for authentication purposes. Additionally, with the new expanded lineup, customers can now start small and grow on the platform with auto scaling. Reduced bandwidth overage fees helps customers deploying bandwidth-intensive applications. These updates allow customers more flexibility and features to deploy their production applications.

Now, turning to our managed-hosting cloud-based offering, we launched malware protection, which detects malware and protects our customers from cyberattacks. This add-on includes critical capabilities such as phishing protection, files protection, database protection for WordPress and Joomla, automated malware cleanup, proactive defense, and cron malware cleanup. These are just a few highlights as we continue to add new capabilities and features to achieve our objective of simplifying cloud and AI infrastructure for our customers. We will continue to listen closely to our customers and accelerate our product philosophy so that customers continue to scale and grow on our platform, which is our primary focus as we work to drive up expansion and improve net dollar retention.

And now, I'll pivot to a part of the business that is seeing a lot of momentum, our AI/ML offerings. We continue to see very strong demand for our AI platform. To support that growing demand and to take the first step of our long-term data center optimization strategy, I'm very excited to announce that we will be opening a new state-of-the-art data center in Atlanta in Q1 of 2025. This not only expands our geographic footprint, providing us cost-effective additional coverage across the US for our core workloads, but also gives us near-term incremental space and power to support our AI strategy and growth.

This new data center is also a key part of our medium-term strategy to reshape our data center footprint, including consolidating workloads from DCs that are currently in expensive locations including New York City, San Francisco, and Toronto, enabling us to improve our gross margin profile over time. As a reminder, opening a new data center gives us ample runway to grow into the additional capacity, and we only add equipment and spend capital as needed to meet demand.

As such, the financial impact of this long-term investment will appear steadily over time as we ramp capacity and leverage it for consolidation of our core workloads, and also for AI training and inferencing as that demand evolves over time. We will share additional details over the next few calls as we start building our new data center out and make further progress on our data center optimization strategy. Given this DC expansion and with the modest increase in AI related capital that Matt will detail in his remarks, it is worthwhile for me to spend a little bit of time providing some context on our AI strategy and how we view this market opportunity.

Today, the majority of AI action across the industry is in the foundational infrastructure layer, with a handful of companies providing GPU infrastructure to a relatively concentrated set of customers that require GPU compute for foundational model training. But over time, we expect generative AI and AI overall to follow a similar progression that the market has seen with other technology evolutions, the action shifting up stack from infrastructure to platforms to eventually applications in the coming years to deliver actual business value to customers.

The heavy users of infrastructure layer today are those building foundational GenAI models, or those extending those foundational models by injecting their own data. This requires a lot of deep expertise in machine learning, data, and foundational models. This restricts AI and associated innovation to well-funded startups and large enterprise companies with very skilled staff given the limited talent pool and high costs associated with this emerging technology, leaving behind the vast majority of companies who don't have access to these capabilities.

Our mission at DigitalOcean is to change this paradigm by democratizing the access to GenAI and AI infrastructure for all customers, just like we did with core cloud computing services using simple-to-use software platform components rather than expensive, CapEx-heavy hardware infrastructure. As a significant step in this direction, we announced the launch of GPU droplets, allowing customers to seamlessly leverage AI technology into their workflows and applications using as few as one or eight GPUs in an on-demand mode.

GPU droplets remove the burden – removes the burden of managing the full lifecycle of GPUs and the orchestration associated with its usage. This type of fractional on-demand access to GPUs is not widely available in the market today. We have seen very robust demand for this capability, which is still in early availability mode.

Additionally, applications that consume AI also need the usual cloud primitives like compute, storage, databases, security, and so on to be deployed in the real world and deliver real business value. Unlike applications that are built on purely GPU farms, software that consumes AI through GPU droplets can seamlessly take advantage of DigitalOcean's core cloud computing platform, making it easy for customers to transition from R&D mode to production very seamlessly, rather than having to go through redeployment.

Let me give you some specific examples of customers that are building on our AI platform. First example is an advanced stage startup building a lightweight but very fast AI code completion tool for developers with a very large context window using native neural network architecture on our platforms. Another example is an AI infrastructure management company that offers a middleware layer to enable rapid training and inferencing for GenAI models on the DigitalOcean platform.

To recap, our longer-term AI mission is more software-centric, with the mission of making it easy for our approximately 638,000 current customers and other companies that look like them, to leverage AI in their application stack without needing super deep AI and machine learning expertise. Now, with Bratin Saha, one of the most accomplished AI leaders in the industry leading the charge for us, we will build on this momentum we have generated over the last couple of quarters and fulfill our mission to democratize AI and make it accessible to all companies.

In conclusion, I'm very pleased with the team's performance in the first half of the year. We have seen growing signs of success in our AI/machine learning business, growth in our core business is reaccelerating, and I'm excited about our near and long-term growth potential across all areas of our business. We have the right leadership team in place and are focused on accelerating our product road map and deliver new capabilities that we announced this year at Deploy, and enhancing our go-to-market motion in the second half of the year.

I will now turn the call over to Matt to provide additional details on our financial results and for our outlook in Q3 and the remainder of the year. Over to you, Matt.

---

## W. Matthew Steinfort

*Chief Financial Officer, DigitalOcean Holdings, Inc.*

Thanks, Paddy. Good afternoon, everyone, and thanks for joining us today. In Q2, we continued to execute on the plans we laid out at the beginning of the year. We made progress on key metrics, we continued to see revenue growth reaccelerate, and we delivered favorable adjusted EBITDA and adjusted free cash flow margins.

Revenue in the second quarter was \$192.5 million, up 13% year-over-year and up 4% quarter-over-quarter. We added \$32 million of annual run rate revenue, or ARR, in the quarter, which was 158% higher than the incremental ARR we generated in Q2 of 2023 and was also the highest step up in nearly two years, excluding the quarter in which we acquired our AI/ML business.



Contributing to this growth was healthy incremental revenue from new customers, increased momentum from our AI/ML platform, which saw significant growth quarter-over-quarter, and contributions from our managed hosting platform, which continues to be one of our faster growing platforms. All of these together are offsetting a flat quarter-over-quarter net dollar retention rate from our existing installed base.

Our Q2 net dollar retention rate was 97%. As we saw last quarter, we continued to see stable performance in net expansion, which is defined as expansion, net of contraction, on our core DigitalOcean platform. Contributing to this stability was our increased product velocity that drove an increase in ARPU, helping to offset the broader macro pressures on net expansion in our customer base.

Our churn levels have also remained very stable for over a year across the business. We are encouraged by the stability in NDR and the modest sequential improvements we are seeing despite the challenging macro environment, which is muting the pace of improvement in NDR and despite a positive but lower contribution to NDR from our managed hosting platform, now that we have fully lapped last April's price increase. We continue to expect stable NDR and expansion levels through the end of the year despite these ongoing headwinds. To further improve our net dollar retention rate, we will continue our solid execution accelerating our product road map, refining our pricing and packaging models, and enhancing our customer success motions.

Beyond NDR, we continue to see acceleration within our AI/ML platform. Q2 ARR has grown over 200% year-over-year from the ARR we acquired last year. We have also successfully navigated much of the initial supply chain and implementation risk that we had identified earlier in the year and are now working aggressively to keep up with demand. We anticipate this momentum to continue for the balance of the year, given the demand for our AI solutions.

Turning to the P&L, gross margin was 61%, which was consistent with the prior quarter and up 100 basis points from the prior year. The 100 basis point year-over-year improvement is primarily a result of the success of our ongoing cost optimization efforts, which to date have more than offset our continued investment in AI infrastructure.

Adjusted EBITDA margin was 42% in the second quarter, which was ahead of guidance and approximately 200 basis points higher than the prior quarter. This beat was primarily driven by strength in gross margin and our ongoing operating cost discipline. Diluted net income per share was \$0.20 and non-GAAP diluted net income per share was \$0.48. GAAP and non-GAAP diluted earnings per share increased by \$0.19 and \$0.04, respectively, on a year-over-year basis. This is a result of our ability to increase our per share profitability levels by driving both operating leverage and reducing our share count. Finally, Q2 adjusted free cash flow was \$37 million, or 19% of revenue.

Turning to our customer metrics, our total Q2 customer count was approximately 638,000, representing an increase from 637,000 customers in Q1. The number of builders and scalers on our platform, those that spend more than \$50 per month, was approximately 161,000, an increase of 7% year-over-year. The revenue growth associated with builders and scalers was 15% year-over-year, ahead of our overall revenue growth rate of 13%. The number of builders and scalers on our platform, which represent 87% of our total revenue, increased by approximately 3,000 quarter-over-quarter.

The continued growth of our larger spending cohorts is a direct result of our focusing our product development and customer success investments on these builders and scalers. The increase in our higher spend and higher

growth customers also resulted in our total average revenue per user, or ARPU, increasing 9% year-over-year to \$99.45.

With our substantial free cash flow generation, our balance sheet remains very strong, as we ended the quarter with \$443 million of cash and cash equivalents. We also continued to execute against our ongoing share repurchase program, and completed \$10 million of repurchases in the quarter.

Moving on to guidance, we expect Q3 revenue to be in the range of \$196 million to \$197 million, representing approximately 11% year-over-year growth at the midpoint of our guidance range. For the third quarter, we expect adjusted EBITDA margins to be in the range of 37% to 38%, and non-GAAP diluted earnings per share to be \$0.39 to \$0.41, based on approximately 102 million to 103 million in weighted average fully diluted shares outstanding.

As a result of the steady performance in our core platform and strong demand we are seeing for our AI platform, we are increasing the bottom end of our full-year revenue guide by \$10 million, projecting revenue to be in the range of \$770 million to \$775 million, a \$5 million increase in the midpoint of our guidance range, and representing year-over-year growth of approximately 11% to 12%. As demonstrated through the first half of 2024, we remain committed to driving continued operating leverage in our core DigitalOcean platform. Given our solid performance in the first half of the year, we are raising our adjusted EBITDA margins guidance for the full year to be in the range of 37% to 39%.

Turning to adjusted free cash flow, we anticipate making appropriate incremental investments through the second half of the year, as we continue to capitalize on the AI opportunity to fuel future growth, although we anticipate these investments having only a modest impact on our cash flow margins. We expect adjusted free cash flow margins for the full year to be in the range of 15% to 17%. As a reminder, adjusted free cash flow can vary quarter to quarter, given the variability of our capital spend and our working capital timing.

From an overall strategic capital allocation perspective, we will continue to be good stewards of our capital and will evaluate opportunities to maximize shareholder return, maintain financial flexibility while continuing to evaluate investments across share repurchases, incremental capacity, and balance sheet management. We are also raising the top end of our prior non-GAAP diluted earnings per share guidance and now expect this to be in the range of \$1.60 to \$1.70.

That concludes our prepared remarks and we'll now open it up 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 Gabriela Borges from Goldman Sachs. Your line is open.

**Gabriela Borges**

*Analyst, Goldman Sachs & Co. LLC*

Q

Good afternoon. Thank you and great to see the step-up in ARR. Paddy, I wanted to follow-up on your comments on AI strategy. Help us understand, are you saying that other GPU providers don't offer the same level of fractional access that you do with Droplets.

And then broadly speaking, comment on the competitive environment. There are a number of well-funded GPU providers that are spending orders of magnitude more CapEx than you are. How do you think about the sustainability of your differentiation when you're up against competitors that have that kind of CapEx spend? Thank you.

**Padmanabhan T. Srinivasan**

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

A

Yeah, thank you, Gabriela, great question. So, as I started the prepared remarks by talking about our AI strategy and how it is differentiated, I want to remind everyone, we are also a very strong player in the GPU Infrastructure as a Service today. So, we have very competitive offerings all the way from bare metal to virtualized environments. And what we announced with the GPU droplets is it's a strategy which enables access to fractional GPU access and also with the ability to orchestrate the lifecycle through the virtualized environment that we provide with GPU droplets. So it just takes the overhead of having to manage the whole infrastructure and makes it super, super easy for those workloads, which we believe are going to be more important for customers that are digital-natives, but also the types of customers that typically reside on our platform.

So they are looking to build applications that consume AI and leverage different parts of the AI infrastructure in a way that extends existing AI models. They are not model builders per se in the sense that they're not trying to build another foundational model. So, we feel our AI strategy, which includes the GPU infrastructure, is tailor-made for customers that are looking to consume AI, not necessarily build foundational models.

So, when I talked about the GPU droplet, that's an abstracted version of the core GPU as a Service. And then we also – in the Deploy conference, we also announced and showed an early preview of what we call as the next generation of our Platform as a Service offerings. So these are endpoint APIs for well-known open-source models like Llama 3.1 and so forth, which we will be releasing in early Q4, which enables another layer of abstraction to build applications on top of LLM platforms.

So we feel our strategy is going more up stack and enabling applications that derive business value from AI, rather than focusing on model builders that are building and training foundational models. So, there's going to be different needs for customers that are looking to derive business value and build applications and platforms on top of our infrastructure. So that's what we are focused on. Certainly, there is going to be room and space for core infrastructure and GPU farms. And there is a very – as I mentioned, very concentrated group of model training companies that are building foundational models, whether it is for question-and-answer, like a large language model for question-and-answer bots, or for text-to-image generation or text-to-video generation and things like that.

But our hypothesis is that that market is very concentrated, and over time, there will be few but very big model companies in the world, and the rest of the universe is going to move towards building applications on top of this infrastructure. And that's where our focus is, and that's historically been our strength is to simplify the access to complex infrastructure by providing the essential building blocks of platform and application building. And that's the strategy we are following.

---

**Gabriela Borges***Analyst, Goldman Sachs & Co. LLC*

Q

Thank you for the detail. Matt, the follow-up is for you on the updated outlook. How did you think about the pace of net new ARR into 3Q and 4Q? Back of the envelope math, I think, implies a moderation from the result that you just put up this quarter. So, help us understand how you thought about handicapping that, call it, \$7 million, \$8 million that you were able to deliver in 2Q into 3Q and 4Q as well.

---

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

A

Yeah. Part of it, as we think about the growth, particularly around the AI side of the business, it's somewhat lumpy. And we talked about earlier in the year that we had some risk baked in. As part of the reason why we gave a wide revenue range in our initial guidance was we had some supply chain risk and that was a new service for us, so we wanted to make sure that we were comfortable with the implementation. And we managed past that risk. And that enabled us to bring on a fair bit of incremental capacity that we had ordered last year in the second quarter, which drove up the ARR. And so, as we look over the balance of the year and we're continuing to add capacity, we're just not – it won't be as lumpy. So it's going to be a little bit lesser than – in terms of the guide than what we had in the second quarter.

The other thing that you're getting when people look at the implied third quarter and fourth quarter growth rates is, from a year-over-year perspective, we've got two lappings that we're doing that are making the growth look like it's decelerating even though we're accelerating clearly in the quarter, and we're showing growth in the other parts of the business. And that's because this is the first quarter that – or third quarter will be the first quarter that – where the Paperspace AI revenue is in the prior period. And so that's a 200 basis point reduction in the growth rate.

And then we're lapping a price increase and customer growth, kind of a nominal customer growth in the Cloudways business that's also contributing about 100 basis points to the year-over-year growth apparent deceleration from second to third quarter.

---

**Gabriela Borges***Analyst, Goldman Sachs & Co. LLC*

Q

Got it. Thank you for the color.

---

**Operator:** Your next question comes from the line of Raimo Lenschow from Barclays. Your line is open.

---

**Raimo Lenschow***Analyst, Barclays Capital, Inc.*

Q

Thank you. Can I go back to the first question that was asked? Just given that we are kind of early on that GenAI journey, so we don't really fully understand how that's going to play out. So, is the message from you guys that as a trainer model, I need like hundreds, thousands, 10,000 GPUs, but as I run a model and if I have an application, I can do that with a much, much smaller GPU number, like the one to eight that you were kind of suggesting today.

Like it just feels a little bit like – because you're early in the journey, it feels a little bit too good to be true. Can you just speak to that a little bit more?

---

**Padmanabhan T. Srinivasan**

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

**A**

Yeah, Raimo, thank you for the question. So, there are three types of AI users today. One is the true foundational model builders, like an Anthropic, Mistral, folks that are building the foundation model. And the second one is what I call it the AI extenders. So they take these models, let's say a Llama 3.1, an open-source model, and they inject or they enhance that using their custom data. It could be a company that has, say for example, a geospatial data, and they enhance an existing LLM with their own custom data source and create a slightly modified version of an existing foundation model. So for this class of customers, they definitely don't need hundreds of thousands of GPUs.

And the third type of user is an AI consumer. So for example, let's say you're creating a new AI-native CRM application or a supply chain application, for which you are relying on a very robust foundational model, again, Llama 3 or Mistral or something. But bulk of your application is to deliver that supply chain forecasting algorithm, but you are leveraging a heavy dose of AI, but you don't need the same kind of physical raw compute power that the category one or category two company needed.

So, these are different types of use cases. So the AI model builders, AI model extenders, and the AI model consumers all have different requirements and need GPU capacity at different scales. And what we are seeing is we are definitely addressing the needs of the second and the third categories that I just talked about.

And another way to think about this, Raimo, is, as with any technology waves, you have infrastructure providers. In this case, this is NVIDIA and all the foundational model builders. They are laying the infrastructure, but the true business value is going to be when this infrastructure is leveraged to build platforms, like simple example would be operating systems based on x86 architecture. And then you have applications which are the ones that truly deliver business value for everyone.

So, as this AI wave goes through – goes up stack from one layer to the other, we feel there is a tremendous amount of need to democratize the access to these GPUs and also provide other software frameworks on the platform layer and infrastructure layer, which is what we are building now.

---

**Raimo Lenschow**

*Analyst, Barclays Capital, Inc.*

**Q**

Okay, perfect. That makes total sense. Thanks for the clarification. And then one for Matt, if I think about the improvement in gross margins you mentioned from just the optimization work you do around the data centers, like what sort of magnitude are you thinking? I don't want guidance, but is this like a few basis points? Is this kind of more meaningful? How should we think about that? Thank you and congrats from me as well.

---

**W. Matthew Steinfort**

*Chief Financial Officer, DigitalOcean Holdings, Inc.*

**A**

Thanks, Raimo. It's a great question and I think I've touched upon this on prior calls. And when I first joined the company, one of the first things that I observed was that the data center footprint was, I'd say, interesting in that we were in tier 1 markets and tier 1 buildings. And that was pretty much the way the company had architected its data center network, which is a relatively expensive way of building your architecture.

Over time, you'd want to move to some tier 2 locations and you'd want to be in more of a wholesale than retail kind of data center kind of footprint. And so, the Atlanta data center that we just announced is a great example of that, where we're going to be able to take down incremental capacity. It will give us room for the expansion of our AI capabilities, but we'll also be able to take workloads out of our more expensive locations and kind of reduce or consolidate those footprints and shift some of the – maybe the less latency-sensitive workloads into different locations.

And that's going to be a multi-year effort. It's not going to be something that we can do over time. But I think there's a meaningful improvement that we can drive in gross margin. And the objective is to continue to drive gross margins up, and that will offset the – in the core business and that will offset some of the lower gross margin mix that we'll get as our AI business continues to grow. So, I'm not going to give guidance on it, but I think it's a meaningful opportunity, and Bratin and I are very focused on driving that [ph] on net (00:40:55) savings and driving the gross margin up.

**Operator:** Your next question comes from the line of Kingsley Crane from Canaccord Security (sic) [Genuity]. Your line is open.

### Kingsley Crane

*Analyst, Canaccord Genuity LLC*

Q

Great. So, first question would be you cited customer success efforts as helping to create stability and churn contraction. Especially now with Larry on board, how much could customer success be used not to just prevent churn but also drive usage in these high potential accounts?

### Padmanabhan T. Srinivasan

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

A

Hello, Kingsley. Good to hear from you. It's a great question. So we currently – our customer success efforts are fairly nascent but ramping up. And they already are tasked at not only managing a relationship with our top customers, but also drive product usage and show them the breadth of our platform, which, after I joined six months ago, we have put some considerable wood behind the arrow and doing two things for our core platform.

One is we are rapidly closing the gap in terms of what our top end of the scalers need from a core product point of view, but also adding a lot of other new capabilities on our App Platform, as I mentioned in the prepared remarks. So, we are doing both, closing the gaps in the product functionality that our top-end scalers need, but also expanding our platform footprint.

So, our customer success reps have – they get trained in all the latest and greatest, and they take that message to our customers and help expand the footprint of these large customers in how they are leveraging our platform. So, that's already there. But I think it is fair to say that it's in its nascency, and with Larry coming on board, this is going to be a big push for us.

So, the purpose of getting someone like Larry is to ensure that he can leverage all his skills in building a high velocity motion to augment our phenomenal product-led growth and make sure that these customers that are getting onboarded to our platform, first of all, know the breadth of our platform, but also start leveraging the capabilities that we have, so that we can drive up expansion.

### Kingsley Crane

*Analyst, Canaccord Genuity LLC*

Q

That makes a lot of sense. And then, I want to return to the GPU topic [ph] of achievements on the (00:43:30) availability. It's great to see that out in the market. And the fractional demand consumption is relatively unique. We've spoken in the past about how software will differentiate you from the other providers. So, just want to check in on Gradient specifically. How important is that to the vision, and how is that acting as an on-ramp for customers?

---

**Padmanabhan T. Srinivasan**

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

A

Yeah. So, we are – so, yes, I completely agree that software as a differentiator is going to be the one that we are going to bank on. And so, right now – so let me again refresh everyone's memory in terms of the offerings we currently have. So, at the bottommost layer, we have a bare metal GPU platform, and we just announced the GPU droplets which provides virtualized and fractional and on-demand access, which is also a pretty big deal. Virtualized, on-demand fractional access to GPUs is what the GPU droplets provide.

And then, a layer up above that is our Platform as a Service offering, which is Gradient, and we are doing a lot of work on that Platform as a Service layer, and we'll be announcing a slew of enhancements over the next 90 days to add a lot of capabilities on the Platform as a Service layer. We believe it is going to be an important front door for customers to start adopting these AI capabilities.

And the abstractions that we are providing right now, like, for example, when customers come into the Gradient platform, they're – yes, they know that they are leveraging GPUs, but that's not why they are coming to our Platform as a Service offering. They're coming because it is super compelling to organize their AI/ML workspaces, and share it amongst their colleagues, and publish it to different repos and things like that. So, there are many different reasons why they come there.

And, oh, by the way, when they have to push something to testing or production, they leverage the GPU infrastructure we have, in which we have a variety of different hardware, not just H100s. We also have a variety of different GPU hardware that helps customers for different use cases. So, I think software is extraordinarily important as we go up the stack, and Gradient is a very differentiated offering for us. And we are working quite a bit on enhancing Gradient over the next 90 days.

---

**Kingsley Crane**

*Analyst, Canaccord Genuity LLC*

Q

Really encouraging. Thank you.

---

**Operator:** Your next question comes from the line of Patrick Walravens from JMP Securities. Your line is open.

---

**Pat Walravens**

*Analyst, Citizens JMP Securities LLC*

Q

Oh, great. Thank you. And first of all, I want to commend you on your recruiting. I mean, the background, like Bratin in particular, is just fabulous. So, someday, I'd like to hear how you did that. But my more specific question is what are the bottlenecks to growing your AI footprint in terms of the data centers? So, data centers that were built for CPUs don't work very well for GPUs, right? So, how are you thinking through that?

---

**Padmanabhan T. Srinivasan**

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

A

Yeah. Pat, thank you so much for the kind words. Yes, you're right, there is surely some constraints. But all of – so, just to be clear, the new data center that I just talked about, we are not live yet on that. So, all of the AI workloads that we are seeing now, we have crammed into our existing data centers, whether it is New York or Amsterdam and other places. Is it easy? No, but we have figured out a way to do that.

See, there are multiple constraints. One is the stuff that you read in the press is power and cooling. Yes, those are really important. The other one is network stack. It's very different from our CPU offering. So, we have to build parallel network stacks to put in these GPUs. And then, real estate, right? The density of our ability to rack and stack GPUs is very different from our traditional density of our racks for CPUs because of the power and the heat that these GPUs throw out.

So, there are some physical constraints, networking constraints, and of course, it is just super expensive and super hard to find more data center space in our existing facilities. So, that's why we decided to go to Atlanta. And by the way, it also helps us consolidate and load balance some of our very expensive data center footprint in the short term.

---

**Pat Walravens**

*Analyst, Citizens JMP Securities LLC*



16 data centers across nine regions that you have in your slide deck. Does that include the Paperspace, whatever Paperspace [ph] had (00:48:30)?

---

**Padmanabhan T. Srinivasan**

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



Yeah, it does. And Paperspace, coincidentally, happen to be in our New York City data center. So, yeah, it includes that.

---

**Pat Walravens**

*Analyst, Citizens JMP Securities LLC*



Okay, great. Thank you.

---

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

---

**Mike Cikos**

*Analyst, Needham & Co. LLC*



Hey, guys. Thanks for taking the questions here. And I just wanted to continue the thought on the DC – sorry, the data center optimization strategy that you guys are working through. With the Atlanta data center that you guys are announcing here, should we think about that as really taking on most of the AI workloads or is it freeing up capacity for you guys to further build out those GPUs in the existing footprint that you have?

And the reason I'm asking is I know when you had gone through the three types of categories, the AI foundational model builders, the extenders, the consumers, I would think that at the very least, the extenders would probably be more concerned about security, i.e. you'd need to be in like a Tier 3, maybe even a Tier 4 data center, given security restrictions. And then can you help us think about the ability to re-architect the capacity that you have for these GPUs? Like I know we're announcing Atlanta today, but should we expect a more significant build-out even beyond Atlanta?

---



**Padmanabhan T. Srinivasan**

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

A

Yeah. So, let me – thanks, Mike. I'll start with the last question. Should we expect a more significant build-out? I don't think so. I think we have enough capacity now for the foreseeable future. But I think the way I look at it, Mike, is slightly differently, which is as we go up the stack, the workload goes from training to what we in the industry call it inferencing. And inferencing, it is important to minimize the latency. So, if you ask me whether all of Atlanta is going to be consumed by AI, no, because that's why Matt and I have been talking about we need to consolidate and load balance across all of our footprint because a lot of the AI capacity will also be in the West Coast or in Canada or in Europe or Asia, because as our inferencing demand grows, some of our customers are already asking us to put GPUs in inferencing mode closest to where their customer density is, whether it is West Coast or Asia.

So, luckily, we have a fairly distributed data center footprint across the globe. So, we are going to take this opportunity to optimize and load balance and consolidate some of these data centers so that we can do couple of things. One, in our core business, we can continue to look at improving our gross margins, get out of these expensive locations, and leverage larger-scale data centers like the one that we just talked about. But, B, as the AI world moves from training to a more inferencing mode, we want to also have the ability to deploy the inferencing GPU capacity closest to where our customer demands are so that we can reduce latency and give them the horsepower they need for inferencing. So, that's how I look at it.

Matt, if there's something you want to add?

**W. Matthew Steinfort**

*Chief Financial Officer, DigitalOcean Holdings, Inc.*

A

No, I think you answered it, Paddy. I think the – to answer the first part of your question, Mike, that it's not going to be filled with – the Atlanta data center is not going to be filled with GPU, but it will take a big chunk of our capacity. And part of that is just getting that GPU capacity that, as Paddy said, we've put in some of our more expensive locations. There's a more cost-effective way for us to do that and we'll migrate a lot of that. That will free up capacity in those more expensive locations for workloads that are more relevant for that, where they require that latency or it's justified to have them in the more expensive locations. So, this is just a great opportunity for us to combine two very important initiatives. One is extending our GPU runway and the capacity we have. And the other is to just load balance and cost optimize in our existing footprints.

**Padmanabhan T. Srinivasan**

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

A

And just to add one more thing, to make it clear, we are not going to rush to fill out the capacity to the brim. So, this gives us land and space and power, but we are going to fill it up as part of our consolidation strategy and as our AI demand grows over the next few years.

**Mike Cikos**

*Analyst, Needham & Co. LLC*

Q

Got it. Got it. And then just a quick follow-up for Matt. Just coming back to the prepared remarks on the net dollar retention. So, is the takeaway here that we're stable at 97% and based on the guide that we have, the teams, assuming that that 97% remains where we are for the rest of the year? And then the second clarification here, but I know last quarter you guys had cited \$90 million in AI-related ARR, can you give us an update on what that is this quarter?

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

A

We're not going to guide – we're not going to provide the details at that kind of platform level at this point going forward. But I think with the 200% increase in – growth at a 200% growth rate, not just increase, but 200% growth rate on what we acquired, you should be able to figure it out in terms of what the ARR was for the AI business. And I'm sorry, the first part of the question was around the 97%...

[indiscernible] (00:54:17)

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

A

...yeah. No, what we're trying to be is, I mean, we're doing a lot to drive NDR up clearly. And our goal is to try to get that above 100% and have that be a tailwind. And we're making good progress on a lot of the initiatives, as Paddy talked about, and the new products and the increased product velocity is contributing. But what we're seeing is just the – primarily the expansion is just stubborn. It's not increasing at the pace at which we'd like to see it. And a lot of that is just our customers aren't growing as fast as they were previously, and their recovery is not as fast as we would like. So, I don't know if we think it's going to be flat. It's just, I mean, we only show it on a rounded basis, and so it looks very flat. But we are seeing kind of green shoots and steady improvement in that.

So, yeah, I think the plan and the guidance that we have doesn't require to get any better than it is today. And so, that, people shouldn't worry about. But we're clearly working aggressively to try to get it up above 100%. And yeah, I'm optimistic that we'll see movement in it, it won't stay completely flat. But it's – because it's such minor movements, we're just kind of saying it's going to be stable.

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

Q

Terrific. Thank you.

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

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

Q

Hello. Great. Thanks for taking the questions and congrats on the results. Paddy, you have really strengthened the leadership team, obviously, with the number of hires that you highlighted. As these leaders kind of settle into their seats, how should we think about some of the changes that they might be trying to drive? And are these changes more focused on for the next year or the second half of this year? And any expenses related to those changes that we should be thinking about in the models?

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

A

Yeah. Thanks, Pinjalim, for the question. I'll answer it in the reverse order. Should we expect to see a different expense profile? No. I think we are going to be super disciplined, and I don't think there's anything fundamentally different that we are going to be doing that will call for a foundationally different-looking P&L or expense profile.

In terms of what we should expect in terms of their impact, I would say, let's take the three key functions, right? On product and technology, Bratin obviously brings a wealth of AI/ML background. So, as we go through the rest

of the year, you should hear a lot from us in terms of both core product innovation, as well as our AI strategy. So, my objective is between now and end of the year, we have to pick up even more pace. As I said, we have doubled the product velocity, but we are looking forward to pushing that even higher and we absolutely want to make our customers super successful on our platform and give them no reason to think about leaving us or leveraging any other infrastructure. So, we want to really ramp up the pace of innovation, and you will see that this year.

The second thing is AI strategy. I want to make sure that we continue to refine our AI strategy in response to our customers. As I've been saying for the last couple of calls, we want the AI strategy to be our strategy to serve our customers. So, you're seeing a lot of different voices in the market in terms of companies that are chasing the big foundational model builders, that's not us. We want to be very disciplined to ensure that we know exactly which customers we're going after and what they need. And I strongly believe that with Bratin here now, we will come out on the right side of this AI strategy and make sure that we are doing it in a very responsible but very expedient fashion.

In terms of Wade, you already saw his team put together a really remarkable Deploy event in just a few days after he joined. We'll continue that [ph] drumbeat (00:58:49). We are bringing in some very strong leaders to increase our presence and our voice in the developer community to drum up interest and really show them that DigitalOcean is back to simplifying their lives, and have us be the real destination choice for developers to learn new complex technologies and start their AI journey from. So, you will start seeing that this year.

Larry, on the other hand, from a CRO point of view, as I said, his mission is to complement our incredible best-of-class product-led growth motion. So, on the sales, on the direct sales and indirect sales, we have a lot of work to do to figure out the best ways to complement the product-led growth motion. On customer success and support, we have a running start. We have a great team and we are going to double-down on the success we already have. And with the product velocity that you are already seeing from us, the customer success team and support team have plenty to work with in taking that message to our customers and help us drive up the NDR over the next several quarters. So, hopefully, that gave you a lot of color, Pinjalim.

---

**Operator:** And that concludes our question-and-answer session. I will now turn the call back over to CEO, Paddy Srinivasan, for closing remarks.

---

### **Padmanabhan T. Srinivasan**

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

Again, thank you, everyone, for spending your hour with us. As I mentioned, we are super excited with the progress. We have a lot of work to do and we are looking forward to staying engaged. And you all have a great rest of the day.

---

**Operator:** 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.