**Datadog, Inc. NasdaqGS:DDOG**

# FQ3 2026 Earnings Call Transcripts

Wednesday, November 4, 2026 8:00 AM ET

# Table of Contents

Call Participants ..... 3
Presentation ..... 4
Question and Answer ..... 9

# Call Participants

## EXECUTIVES

**Alexander Joseph Ljubich**
*SVP of Strategic Finance & Investor Relations*

**David M. Obstler**
*Chief Financial Officer*

**Olivier Pomel**
*Co-Founder, CEO & Director*

## ANALYSTS

**Gabriela Borges**
*Goldman Sachs Group, Inc., Research Division*

**Samik Chatterjee**
*JPMorgan Chase & Co, Research Division*

**Patrick Edwin Ronald Colville**
*Scotiabank Global Banking and Markets, Research Division*

**Matthew George Hedberg**
*RBC Capital Markets, Research Division*

**Koji Ikeda**
*BofA Securities, Research Division*

**Karl Emil Keirstead**
*UBS Investment Bank, Research Division*

**Raimo Lenschow**
*Barclays Bank PLC, Research Division*

**Stewart Kirk Materne**
*Evercore ISI Institutional Equities, Research Division*

**Joel Meshack Omino**
*Citigroup Inc., Research Division*

**Sanjit Kumar Singh**
*Morgan Stanley, Research Division*

**Peter Weed**
*Bernstein Institutional Services LLC, Research Division*

**Aleksandr J. Zukin**
*Wolfe Research, LLC*

# Presentation

**Operator**
Good day, and thank you for standing by. Welcome to the Third Quarter 2026 Datadog Earnings Conference Call. [Operator Instructions] Please be advised that today's call is being recorded. I will now hand it over to our first speaker, AJ Ljubich, Senior Vice President, Strategic Finance and Investor Relations. Please go ahead.

**Alexander Joseph Ljubich**
*SVP of Strategic Finance & Investor Relations*

Thank you, Victor. Good morning, and thank you for joining us to review Datadog's third quarter 2026 financial results, which we announced in our press release issued this morning. Joining me on the call today are Olivier Pomel, Datadog's Co-Founder and CEO; and David Obstler, Datadog's CFO. First, I'd like to thank Yuka Broderick for her excellent work leading our IR function over the last 5.5 years as she takes on a new role within Datadog.

I'm excited to step back into the IR function after having worked with many of you through the company's 2019 IPO. During this call, we will make forward-looking statements, including statements related to our future financial performance, our outlook for the fourth quarter and for the fiscal year of 2026 and related notes and assumptions, our product capabilities and our ability to capitalize on market opportunities.

The words anticipate, believe, continue, estimate, expect, intend, and will and similar expressions are intended to identify forward-looking statements or similar indications of future expectations. These statements reflect our views as of today and are subject to a variety of risks and uncertainties that could cause actual results to differ materially. For a discussion of the material risks and other important factors that could affect our actual results, please refer to our Form 10-Q for the quarter ended June 30, 2026.

Additional information will be made available in our upcoming Form 10-Q for the fiscal quarter ended September 30, 2026, and other filings with the SEC. This information is also available on the Investor Relations section of our website, along with a replay of this call. We will discuss non-GAAP financial measures, which are reconciled to their most directly comparable GAAP financial measures in our earnings release, which are available at investors.datadoghq.com. With that, I'd like to turn the call over to Olivier.

**Olivier Pomel**
*Co-Founder, CEO & Director*

Thanks, AJ, and thank you all for joining us today to go over our Q3 results. Let me begin with this quarter's business drivers. I am pleased to say that we executed a strong Q3. We saw strength across our business, including across our product lines and across our diverse customer base. Revenue growth from our non-AI customers accelerated for the fifth straight quarter and was in the high 20s percent year-over-year, up about 1 percentage point from last quarter and up from about 20% in the year-ago quarter.

Additionally, our AI customer cohort remains a strong growth driver. This cohort continued to grow and diversify, both in the number of customers we serve and in the scale of those customers. Overall, we continue to see healthy demand as customers from the most tech-forward startups to the largest enterprises in the world are building with AI and with modern technologies, which we believe is fueling expanded adoption of the cloud and of the Datadog platform.

Regarding our Q3 financial performance and key metrics. Revenue was $1.2 billion, an increase of 35% year-over-year and above the high end of our guidance range. We ended Q3 with about 33,500 customers, up from about 32,000 a year ago. We also ended with about 4,860 customers with an ARR of $100,000 or more, up from about 4,060 a year ago. These customers generated about 91% of our ARR.

And we generated free cash flow of $251 million with a free cash flow margin of 21%. Turning to product adoption. Our platform strategy continues to resonate in the market. 61% of our customers now use 4 or more products, up from 54% a year ago. 41% of our customers use 6 or more products, up from 31% a year ago and 16% of our customers use 10 or more products, up from 8% a year ago. So we're landing more customers and delivering value across more products through our unified platform. One milestone we are very proud of, our Infrastructure Monitoring suite of products now exceeds $2 billion in ARR, growing in the mid-30s percent year-over-year.

Infrastructure Monitoring was the first product we brought to market more than 10 years ago. Today, our customers rely on Datadog for visibility across their environments, including cloud infrastructure, on-prem servers, containers, serverless applications and GPU fleets. And while we're thrilled to reach this milestone, we still feel that we're just getting started. According to industry analysts, we are the market share leader in global observability, but with a percentage market share still in the teens.

Moving on to R&D. As AI changes how our customers build and run software, we've been investing and executing on 2 fronts: AI for Datadog and Datadog for AI. First, AI for Datadog. These are AI products and capabilities that make the Datadog platform more useful and increasingly automated for our customers. First off, our flagship agent, Bits AI, is off to a very strong start. Since we launched Bits, it has become the fastest-adopted product in our history. And in October, Bits Testing became generally available.

With a single click, it discovers the critical user flows in an application and generates tests to cover them. Meanwhile, our AI Research team is hard at work pushing the boundaries of what Bits and our platform can do. Last month, we introduced Scout, our own reasoning model specifically optimized for SRE tasks.

Scout was built using our in-house distributed reinforcement learning infrastructure and now powers our Bits Investigations. This is exciting for 2 reasons. First, it performs on par with or even better than frontier models on alert investigations at a fraction of the inference cost.

This is valuable as it allows Bits to investigate many more alerts, start earlier and go deeper on incidents as an always-on first responder for our customers. Second, this is also exciting as a demonstration of what we can achieve when we train models with our own production data, critical user workflows and internal tooling. Scout is only the beginning. With the Adaptive ML team, which joined us in June, we are broadly accelerating our work on reinforcement learning and model post-training.

Over time, we expect our research efforts to drive a step change in capabilities across our platform and move us closer to our goal of autonomous operations. Moving on next to Datadog for AI. These are the products that deliver end-to-end observability and security across the new AI stack. As AI moves from experimentation into production, customers today need to answer 3 key questions, how is it performing? Is it secure? And what's the ROI or is it worth the spend?

And because AI doesn't run in isolation, they want these answers on the same platform as the rest of their tech stack. On performance, we cover the whole AI stack from compute all the way to agents. GPU Monitoring now goes beyond hardware reliability with new capabilities that help customers maximize fleet utilization, increase training throughput and serve inference at a lower cost per request. The number of GPU devices we monitor is up by a factor of 8 so far this year, and some of our frontier lab customers are working with us as design partners on these capabilities.

Going up the stack, Agent Observability helps teams monitor and evaluate their agents in production, including the quality of their responses, latency and token usage. The number of spans sent to Agent Observability is now up by a factor of 9 this year. We've also made Agent Observability available in the Anthropic marketplace, and were a launch partner to OpenAI's new enterprise marketplace as well. So now customers can buy our product directly through the frontier labs they already work with.

On the security side, our AI Guard product discovers the custom agents running across an organization and protects them at runtime against prompt injection and data exfiltration, drawing on the same real-time data we collect to observe them. And on the cost side, companies everywhere are grappling with a very fast growth in token and AI spend. Beyond the 3 major clouds, AI providers are now the largest category of spend our customers monitor in Datadog Cloud Cost Management, or CCM.

Now with AI Costs in CCM, customers can see which teams, models and agents drive their AI spend, and can govern that spend with budgets, forecasts and anomaly alerts. More than 1/3 of our CCM customers now track AI costs. The number of customers doing so is up more than 11x. Moving up the stack again. Our Agent Console adds to our CCM data and shows exactly which agents are actually used by whom, for what tasks, at what cost and most importantly, with what outcomes.

As we continue to innovate across AI for Datadog and Datadog for AI, independent research is taking note. Forrester named Datadog a leader in its 2026 Forrester Wave for AIOps Platforms. And Gartner's August 2026 AI Vendor Race report called Datadog the "Company to Beat" for observability platforms. We believe this validates our leadership in observability as well as our ability to help customers manage complexity in the age of AI.

Let's move on to sales and marketing and look at a few of the deals our go-to-market teams have closed in what has been a very strong quarter. First, we signed a 7-figure annualized expansion with a leading software company, more than doubling their annual commitment. This customer is moving AI agents into production and will replace a stand-alone commercial tool with Datadog Agent Observability to monitor their cost and quality. And by using Flex Logs, they can gather and correlate larger log volumes at a lower cost.

In total, they will consolidate 3 tools onto Datadog and grow to 20 Datadog products, including On-Call and Bits AI. Next, we landed a 5-year 7-figure annualized deal with a U.S. federal defense agency. This customer runs thousands of cloud accounts across hundreds of mission owners. And as they've modernized, their manual tooling left them with fragmented visibility and security blind spots. They chose Datadog as their enterprise observability and security platform and will start with 9 products, including Infrastructure Monitoring and Cloud Security.

I want to congratulate our public sector team on this milestone win, which was a direct result of our investment in government and the FedRAMP High certification we secured earlier in the year. And we are actively expanding our government platform to support workloads up to Impact Level 5, or IL5, the highest level of accreditation for unclassified data, which is part of our long-term investment in the U.S. public sector. Next, we landed a 7-figure annualized deal with a leading home healthcare software company.

As this customer has grown, fragmented tools have slowed incident response and hurt customer experience. They will consolidate 4 tools into Datadog, replacing their legacy log and security vendor with Flex Logs and Cloud SIEM and relying on our Sensitive Data Scanner to help them stay HIPAA compliant. This customer will start with 9 Datadog products, including Incident Management and Bits AI for faster resolution.

Next, we signed a 7-figure annualized expansion with a leading online entertainment company. This customer is seeing a surge in traffic with their peak season and is scaling effectively with Datadog. Their own AI agents already use our MCP Server, and today, they are also adding Bits AI. This customer will grow to 23 Datadog products, including Agent Observability and Cloud Cost Management. Next, we landed 7-figure annualized deals with 2 of Japan's most important banks. As they modernize their core financial systems, both banks are standardizing on Datadog for unified visibility across cloud and on-prem environments.

Both customers are adopting the full Datadog platform, including all 3 pillars: Digital Experience, Security, Software Delivery, Service Management and Bits AI as they move towards AI-driven operations. Next, we landed a 6-figure annualized deal with a top European AI lab. This customer's developers were losing valuable time to slow builds and flaky tests, which delayed releases in a highly competitive market.

They chose Datadog CI Visibility and Test Optimization, replacing a commercial tool and freeing their own engineers to focus on shipping products. Their engineers also use Bits AI to quickly build monitors and custom analytics. Next, we signed a 7-figure annualized expansion for an 8-figure annualized deal with a leading AI coding company. This customer is growing rapidly and renewed early to keep pace with their expanding footprint.

With this expansion, they will adopt Flex Logs and our new Infinite Cardinality Metrics to scale cost-effectively with volume. And they are also one of the largest users of our MCP Server. This customer will

grow to 13 Datadog products. Finally, we signed an 8-figure annualized expansion with the AI division of one of the world's largest technology companies, which nearly triples their annual commitment. Just 2 quarters after we landed them, they are scaling Datadog across a large number of AI training workloads. They are also adopting APM and Bits AI to further reduce engineering friction.

Alongside these wins, on the other side of the customer spectrum, we are leaning into the large number of new AI-native companies we see created every day by investing in our Datadog for Startups program. We are broadening its reach to a larger number of start-ups that are earlier in their lifecycle and offering deeper engineering collaboration with the ones that are seeing explosive growth. This is a truly exciting time to partner with start-ups. Now before I turn it over to David for a financial review, I'd like to say a few words on the longer-term outlook.

There is no change to our overall view that digital transformation and cloud migration are long-term secular growth drivers for our business. And AI is another large growth driver for us today. As AI workloads move into production, they drive more cloud consumption and more usage of the Datadog platform, and we see that today across our business. All 15 of the top 15 AI leaders are our customers. And more broadly, we now serve over 850 AI customers, up from about 750 last quarter. And across our broader customer base, we also see AI activity growing substantially.

About 1/3 of our customers now use one or more of our AI products, including 3/4 of our $100,000-plus customers with many of these customers still just getting started with these products, leaving a large opportunity for future expansion. And we continue to see rapid growth in agentic activity with the number of MCP tool calls up by a factor of 50 so far this year. We are tremendously excited about where this is going. As AI writes and runs more of our customers' software, they're shipping more code faster across a growing mix of models, agents and infrastructure.

Of course, this generates complexity that is very quickly outpacing what even the best engineering teams can track on their own. And that's exactly the problem we're investing to solve. AI for Datadog takes on more of that work for them. Datadog for AI helps them run their AI with confidence and our AI research turns our data into models that will deliver more autonomy over time. Because of that, now more than ever, we feel ideally positioned to help customers of every size and every industry as well as every type of user, whether humans or AI agents, so they can transform, innovate and drive value through AI and cloud adoption. And with that, I will turn it over to our CFO. David?

**David M. Obstler**
*Chief Financial Officer*

Thanks, Olivier, and good morning to everybody. Our Q3 revenue was $1.197 billion, up 35% year-over-year. We continued to see robust usage growth from existing customers as well as a strong ramp in our new customers. Revenue growth accelerated within our broad base of customers, excluding AI customers, to the high 20s percent year-over-year, up about 1 percentage point from last quarter and up from about 20% in the year-ago quarter.

And we saw robust growth with broad-based strength across customer segments, spending bands and industries. And excluding our largest customer, revenue growth accelerated again. As we communicated in August, our largest customer reduced their usage in the third quarter, while their usage remained above their minimum commitment. Meanwhile, our AI customers' growth continues to significantly outpace the rest of the business. This over 850-strong group continues to diversify and grow, including 39 customers spending more than $1 million annually with us. And we continue to land larger and grow faster with new logos, contributing more to revenue growth.

The portion of our year-over-year revenue growth that relates to new customers was about 30% in Q3, up from 25% in the year-ago quarter. And now regarding our retention metrics. Our trailing 12-month net revenue retention percentage was in the low 120s percent similar to last quarter, and churn remains low, with gross retention stable -- gross revenue retention stable in the mid-to-high 90s. We believe this metric highlights the mission-critical nature of our platform for our customers.

And now moving on to our financial results. Billings in the quarter were $1.16 billion, up 30% year-over-year. Remaining performance obligations, or RPO, was $3.85 billion, up 38% year-over-year. And current RPO grew in the mid-30s percent year-over-year. RPO duration increased slightly year-over-year. As a reminder, we continue to believe revenue is a better indicator of our business trends than billings and RPO. Let's review some of the key income statement results. Unless otherwise noted, all metrics are non-GAAP.

We have provided a reconciliation of GAAP to non-GAAP financials in our earnings release. For Q3, our gross profit was $972 million, and the gross margin was 81.2%. This compares to a gross margin of 79.6% last quarter and 81.2% in the year-ago quarter. Q3 OpEx grew 25% year-over-year versus 26% last quarter and 32% in the year-ago quarter. And Q3 operating income was $334 million or 28% operating margin compared to 23% last quarter and 23% in the year-ago quarter.

And now turning to the balance sheet and cash flow statements. We ended the quarter with $5.2 billion in cash, cash equivalents and marketable securities. Cash flow from operations was $290 million in the quarter. After taking into consideration capital expenditures and capitalized software, free cash flow was $251 million with a free cash flow margin of 21%. Quarterly, the timing of collections and payments can vary and on a 12 -- trailing 12-month basis, free cash flow margin was 26%. And now for our outlook for the fourth quarter and for fiscal year 2026. First, our guidance philosophy overall remains unchanged.

As a reminder, we base our guidance on trends observed in recent months and apply conservatism on these growth trends. Regarding our largest customer, our Q4 guidance contemplates their revenue at their minimum commitment. Now for the fourth quarter, we expect revenue to be in the range of $1.19 billion to $1.2 billion, which represents a 25% to 26% year-over-year growth. Excluding our largest customer, Q4 revenue growth implied in our guidance is in the low to mid-30s percent range.

Non-GAAP operating income is expected to be in the range of $280 million to $290 million, which implies an operating margin of 24%. The fourth quarter includes large marketing investments, including AWS re:Invent. Our non-GAAP net income per share is expected to be in the $0.69 to $0.71 per share range based on approximately 379 million weighted average diluted shares outstanding. For the fiscal year 2026, we expect revenue to be in the range of $4.515 billion to $4.525 billion, which represents approximately 32% year-over-year growth.

And non-GAAP operating income is expected to be in the range of $1.095 billion to $1.105 billion, which implies an operating margin of 24%. Finally, non-GAAP net income per share is expected to be in the range of $2.73 to $2.75 per share based on approximately 375 million weighted average diluted shares outstanding. Now some additional notes on our guidance. We expect net interest and other income for the fiscal year 2026 to be approximately $200 million.

We expect cash taxes to be about $30 million to $40 million, and we continue to apply a 21% non-GAAP tax rate for 2026 and going forward. And finally, we expect capital expenditures and capitalized software together to be in the 3% to 4% of revenue range in fiscal year 2026. To summarize, we are pleased with our execution in Q3, which combined strong growth with strong profitability. We continue to invest to capture our long-term opportunity in the cloud and in AI. And I want to thank Datadogs worldwide for their efforts. Now with that, we will open the call for questions. Operator, let's begin the Q&A.

# Question and Answer

**Operator**
[Operator Instructions] Our first question will come from the line of Sanjit Singh from Morgan Stanley.

**Sanjit Kumar Singh**
*Morgan Stanley, Research Division*

Congrats on the strong results on Q3. David, I want to start with the Q4 guide. If we wind back 90 days ago, has anything changed with respect to your underlying assumptions between the different cohorts, so between the core business and the sort of largest customers versus 90 days ago? Or is this outlook just sort of a reflection of that sort of minimum commitment methodology that you're applying in Q4? Just want to get a sense of has anything changed with respect to the outlook...

**David M. Obstler**
*Chief Financial Officer*

Yes, there are a couple -- yes, thank you, Sanjit. There are a couple of things to talk about. First of all, I think we said versus our largest customer that they renewed and they were spending above the minimum commit, and that happened in Q3. We said that in our guidance, we were going to only include the minimum commit despite the fact that in Q3, they were spending above the minimum commit. The -- as we've talked about, we've accelerated in Q3, excluding the largest customer.

So I would say the only thing that has changed was in sort of Q3, we did some renewals extensions with some large and large AI customers where as normal, we essentially extend and increase the commit, but do better pricing for that. And that happened in a number of large customers. It's a normal behavior, and that's really the only thing that changed since our last communication. Oli, do you want to add anything?

**Olivier Pomel**
*Co-Founder, CEO & Director*

Yes. There's a few -- we discussed on the call, I could bring up a few of those large customers where we saw renewals way earlier than we thought we would. There's another one that happened in October that we're not discussing. So these are all good news, but they do create like headwind into the short term into the revenue, and that's one change from last quarter. Other than that, like everything has played out pretty much the way we were expecting with some happy surprises increasing, in particular, the continued acceleration of the non-AI customer, which has been very good overall.

**Sanjit Kumar Singh**
*Morgan Stanley, Research Division*

Understood. And if we think about the ex-AI native customer growth, that's been on an accelerating trend for several quarters now. As we think about this core non-AI native business, you obviously have a lot of tailwinds on the AI portfolio. It seems like the $100,000 customer quarter is more of the AI product. Do you feel that, that can continue to drive sort of sustained trajectory? Or are we -- after multiple quarters of acceleration, do you sort of see a leveling off in that ex-AI native cohort?

**Olivier Pomel**
*Co-Founder, CEO & Director*

I mean, look, in the [indiscernible] that all of these customers are adopting AI and they are continuing more cloud and they're building more software and they are changing the way they work. So we expect a long period of growth for these customers. Now with the overall growth rate is, does it keep going up a lot, does it stay stable? It's a very high growth rate already, and we feel pretty good about it. It's hard to tell where it's going to go. But we're very confident in the broader trend that is underpinning that.

**Operator**
Next question will come from the line of Gabriela Borges from Goldman Sachs.

**Gabriela Borges**
*Goldman Sachs Group, Inc., Research Division*

David, I want to stay on this point on the AI native cohort because I think you've been very transparent in how you talk about your largest customers versus the rest of the cohort. And my question for you is you have these funky sequentials because you have had this incredible growth rate in this cohort over the last 12 months. And now you've got this renewal dynamic with your largest customer, you're also talking about the renewal dynamics on the early renewals.

So my question for you is, what can you tell us about how to think about modeling the sequentials in this business? Is it that we get through the next, call it, 2 quarters and then we're back to steady state? Is it that maybe we see another step function down every year in the sequential because you have this early renewal dynamic, which again, is a really good place to be because you're getting these really large deals but the only way you get better pricing is because the dollar volume is going up. So maybe just the puts and takes to how we think about the sequentials for the AI native cohort and when we will be back to something that looks like normal, which I get is a tricky definition of the word normal anyway.

**David M. Obstler**
*Chief Financial Officer*

Yes. So good question. So in terms of the largest customer, I think as we said, the customer renewed -- the customer -- we decided to put only the commitment, the minimum commitment. So the customer in Q3 spent substantially above the minimum commitment. So one thing that is very important in the sequentials is what happens there. We don't control that. The customer can't go below the minimum commitment, but can be various degrees above the minimum commitment. So that is one factor that we -- so essentially, one thing that will create less of a sequential guidance or actual would be that, that number stabilizes and isn't running above the minimum commitment. So that's one.

**Olivier Pomel**
*Co-Founder, CEO & Director*

I would say and I will broaden it, I mean, the point here is there's one part that is very usual, which is that we have all these customers that are expanding and they're growing and they're renewing. And then when they renew, they get better deals and they also adjust their traffic over time with an underpinning of growth, which is what we've seen since the beginning of the company. There's something that is new that we have, some of these companies are scaling incredibly fast to a level that we've never seen in the industry or any industry really before, which creates new dynamics.

This is why we can be a bit surprised when we see customers renew early just a few quarters into there -- starting to add up the product. In general, though, we're getting a larger number of these customers like the AI native cohort is broadening. We have many more customers, many more large customers into that cohort. We called out a few wins that we landed another AI lab, for example, this quarter that's in the note.

So we -- the long term, this is going to broaden. This is going to represent a much larger fraction of our customer base. And as a result, there will be more stability over time. I still don't think it's a very unstable business. I mean we're talking about a change of a few percentage quarter-to-quarter in terms of where we end up versus where we think we're going to be. But we think at the end of the day, this is a usage-based model, and we get more visibility into the mid- to long term than we do into the very short term.

**David M. Obstler**
*Chief Financial Officer*

I do want to remind everybody that all things being equal, the sequential between Q3 and Q4 tends to be less because of the late December effect when our usage-based products and vacations, so -- and some

code freezes. So that can be transparent based on the other factors or it might be underneath of that, but we do have that effect in Q4, which we don't have in the other quarters.

**Operator**
Our next question will come from the line of Raimo Lenschow from Barclays.

**Raimo Lenschow**
*Barclays Bank PLC, Research Division*

Can I stay on one last time -- one last question on this one. So if I think about the customers renewed early. So they get better pricing. But obviously, the volume is going -- kind of going crazy there, but it kind of creates a short-term headwind because the same volume is now on a lower price, but then the volume grows and from then on in a couple of quarters, we should see like the normal progression. Is that kind of the right thing to think about? I think we have that in the past, but I just wanted to clarify that. And then I have one follow-up.

**Olivier Pomel**
*Co-Founder, CEO & Director*

Yes, that's broadly what happened. There's a few other things that typically happen in those situations, which is they try to understand what is that going to need. So they look at what they're using and they think, okay, this part, maybe we don't need that much of -- this part, we're probably more of or maybe in some situations, we tell them, hey, maybe you should use this different tier of retention for this type of data because of the way you're using it.

So there's a bit of account optimization there when we work with them. But the -- that this is on the background of growth. These customers are growing and they're deploying much larger workloads, and they're looking at doubling, tripling, quadrupling their footprint with us over time, and this is basically we're working on with them.

**Raimo Lenschow**
*Barclays Bank PLC, Research Division*

Okay. Perfect. And then the main question for me was on the ex-AI part, like historically, we would think about that growth coming from cloud migration, people modernize and doing more stuff. You mentioned another aspect, which is kind of because of AI, et cetera, we can code more, there's more application modernization, all these stuff happening. So is this kind of changing the way it's kind of evolving for you from classically with people modernize, you have to go to Datadog. Thank you very much. Well, now actually the companies are working differently, and that's also starting to drive the ex-AI business. Can you speak to that, please?

**Olivier Pomel**
*Co-Founder, CEO & Director*

Yes. I mean -- well, the one thing I didn't mention on the call, that's always true. We also do more things for our customers. We have consolidated a lot of what they were using. We have more products to cover more of the surface area. As it comes to customers building more with AI, like it also brings some of our other products more into the front and center. So one example is customers are building a lot more software. So they're putting a lot more pressure on their CI and deployment and testing. And so they're trying also to adopt more of our CI products. I think we called out a few wins around that too in the call. So that opens the door for more products beyond core observability for us.

**Operator**
Our next question will come from the line of Samik Chatterjee from JPMorgan.

**Samik Chatterjee**
*JPMorgan Chase & Co, Research Division*

I had a couple. Maybe for the first one, I know you mentioned your largest customer did spend above their commit in Q3. Is there any way to tell us how much of the beat relative to your guidance in Q3 was because of the large customer spending above the commit relative to the rest of the business, right? Depending on how you forecasted them in Q3? And I have a quick follow-up.

**David M. Obstler**
*Chief Financial Officer*

Yes. We're not answering -- we're not being specific, but it is true. It just -- it is true that if you're basically guiding like in below there -- we said we'd be conservative. In Q3, we said conservative relative to what they were spending. And in Q4, we've been explicit and said minimum commitment. It is true that if they spend more than that, that would be implicit beat mathematically, yes, yes.

**Olivier Pomel**
*Co-Founder, CEO & Director*

Again, we didn't change the assumption there over time. So we get everything unchanged. I think the one difference -- we brought up earlier, the one difference since the last quarter is some positive trends on the overall customer growth, but also some early renewals of larger customers, which includes some other AI customers. That's the one thing we should remember as being the difference.

**Samik Chatterjee**
*JPMorgan Chase & Co, Research Division*

Yes, sure. And for my follow-up, I mean, the contribution of revenue from new customers continues to ramp up steadily. Can you just talk about sort of what is the average land size that you're seeing now with new customers and how that's changed? And what sort of behavior you're seeing on that front? How you're thinking about it as you sort of move into next year in terms of how much new customers contribute to your overall growth?

**David M. Obstler**
*Chief Financial Officer*

We haven't provided guidance on -- we said over time, we've seen an increase of average land size due to the increase of the platform and as we give the metrics on the adoption of products, the adoption of more parts of the platform. Oli, do you want to...

**Olivier Pomel**
*Co-Founder, CEO & Director*

Yes. I think -- look, this is something we see in the customer base that we see that one is -- one, with the new and the AI native companies, some of them start with more scale earlier, like we've discussed it before, but we see companies scaling faster than used to. So they have -- they represent more of the revenue faster. But we also are developing new segments and landing large customers from day 1.

So we called out a few on the script today. We mentioned, I think, a defense agency in the U.S. government. We mentioned Japanese banks, like 2 of the 3 largest Japanese banks in the same quarters. Those are all sizable low land deals. And if you compare the business to what it was, say, 5 years ago, like this is definitely a new side of business that we didn't necessarily have before. And by the way, those wins I mentioned, they don't even register into the revenue yet because these are lands but these are not customers that have ramped usage yet. So...

**Operator**
Our next question comes from the line of Matt Hedberg from RBC.

**Matthew George Hedberg**
*RBC Capital Markets, Research Division*

I wanted to start with training. You noted in your script, you had an 8-figure, I think, annualized expansion with an AI division of a large software company. It sounds like they're using you guys for AI training workloads. I guess this is -- to us, it seems like a really interesting growth opportunity. Can you step back and kind of talk about what that training opportunity is? And what could it mean for your growth?

**Olivier Pomel**
*Co-Founder, CEO & Director*

I mean it's been evolving positively. If you look a year or 2 years ago, basically nobody was doing training. It was only a few frontier labs who were doing their own training and nobody else was. Now we see an order of magnitude more large-ish labs or new labs that are doing training at a large scale. And then there's another whole category of companies below that, that are starting to do their own, including us. We started doing training at a pretty decent scale, and we showed some of the results, we showed some of the results on the call as well.

So I think there's some interesting potential development there where training becomes a part of everybody's business and everybody's customizing models and building their own. And so we are investing. I would say it's still early to call it as a mass market. Right now, we see some interesting indications of it. We also see that we can be extremely useful to some very large customers. But I think we still have to invest. We have to build. We have to see how that market develops.

**Matthew George Hedberg**
*RBC Capital Markets, Research Division*

Got it. And actually, David, if I could just have one more follow-up. On your largest customer, you used the word substantially above their minimums, which I find interesting. And I know you're assuming only the minimums for 4Q, but philosophically, looking at this customer and their recent consumption patterns, does this perhaps change your opinion about how they might use Datadog longer term, maybe new use cases above and beyond what they're using today that could make that growth opportunity really interesting for this customer?

**David M. Obstler**
*Chief Financial Officer*

I think that's true for any customer that there are many different use cases. For that, in particular, we -- as we've said all along for many quarters, we don't control. We have the contract, the renewal, and we're not in the business of projecting that. But it is true from -- across our customer base that there are many use cases in a company, whether it be division and end user, et cetera. Oli, you want to...

**Operator**
Our next question will come from the line of Karl Keirstead from UBS.

**Karl Emil Keirstead**
*UBS Investment Bank, Research Division*

Okay. Great. Maybe, Olivier, for you. You mentioned in your comments that a lot of customers are grappling with higher token costs, higher AI spend and are trying to get it under control. I'm just wondering, based on all your customer interactions, how enterprise spending behavior is changing in light of the fact that they're trying to get their token costs under control. And to the extent that any of the customers that early renewed to try to get price discounts were not AI natives. They were traditional enterprises. Do you think that might be because they're trying to grapple with budget pressure?

**Olivier Pomel**
*Co-Founder, CEO & Director*

Yes. So I think, look, this is a very widespread problem. I think what we've seen is there's been an explosion of adoption of AI, in particular, coding agents. And that was not planned largely. A year ago,

nobody had the amount they spend in their budget on AI. So they really had to try and figure out what that works and what part of that they actually need and what they get from that. So there's a strong need for tooling, which is something we provide.

It's a great opportunity for us, and we see great growth and adoption of our Cloud Cost Management products on that or Agent Console that is our products that -- now in terms of the early renewals that we see, they're not directly related to that. I think they come from large customers that are either AI labs or that are focused on -- that are tech companies in general, I would say. And they -- and that doesn't really have anything to do with their own -- they think something has more to do with their own growth as a business.

**Operator**
Next question will come from the line of Fatima Boolani from Citi.

**Joel Meshack Omino**
*Citigroup Inc., Research Division*

This is Joel on for Fatima. I'll maybe just ask 2 real quick. So firstly, just on margins and operating and gross income leverage, maybe for you, David. Strong results this quarter, at least relative to our expectations, but we do know there's also certain areas that you are investing fairly heavily in. So just curious what some of the drivers of the performance were specifically this quarter? And then what your updated or what your latest investment posture is? And then just the quick second question, I'll just throw it in now is any data points on Bits AI penetration specifically in the customer base?

**Olivier Pomel**
*Co-Founder, CEO & Director*

Do you want to go on margins first?

**David M. Obstler**
*Chief Financial Officer*

Yes, I'll go with margins first, yes. Yes, it's very similar to what we said. We are investing aggressively in the go-to-market, we continue to have white space uncovered, both from a geographical or an end market, for instance, Fed. And we are continuing to be successful in scaling our quota capacity and you see that in the growth.

In addition, we're continuing to invest in R&D. In this case, it's a combination of AI, of investments in a combination of AI and people. We see lots of opportunities. I think Oli will now go -- talk about the products, but we continue to see a lot of opportunity for us to get market share, get platform penetration and increase our feet on the street or on a global basis. Want to go over the Bits?

**Olivier Pomel**
*Co-Founder, CEO & Director*

Yes. And on Bits, I think we've mentioned on the call, we see amazing adoption of it, like this is -- when you look at the number of individual customers or users that are using it, this has been the strongest launch we've ever had as a product. So we feel very good about it. We've also revamped the way we package it over the past quarter. So we used to charge specifically for investigations.

Now we have a model of AI credits that cover a whole different set of use cases that involve Bits as our agent and we've seen great results. Thanks to that. So it's super encouraging. I would say it's still early because there's still so much that we were building into this. So many things our customers want us to do for them, so many opportunities to automate their workflows. But we are -- we feel very, very good about the current status.

**David M. Obstler**

*Chief Financial Officer*

Let me just circle back on one thing. I think I neglected to go over gross margin, as you mentioned. I think it's very similar to what we said before that through a combination of the penetration of our platform as well as our engineering, we have been plus or minus in this 80% range. I think you see in this quarter, we ended up around 81%. And so this is common practice for us. We invest, optimize, do engineering projects and what was communicating gross margin indicates stability there, which we've evidenced in previous quarters.

**Operator**
Our next question will come from the line of Kirk Materne from Evercore ISI.

**Stewart Kirk Materne**
*Evercore ISI Institutional Equities, Research Division*

I guess, David, on -- or Oli, on this native AI cohort that's renewing a little bit early. Obviously, they're coming back because they're spending so much with you all and obviously want to get better economics. I'm just curious, what's the sort of attach rate on some of those newer renewals? When they come back and renew early, do you still have time to sort of cross-sell them some of the other products? I guess what are you seeing in that dynamic? I realize it creates some near-term headwinds from a revenue perspective, but it would also, I think, open up the opportunity to have a discussion about you expanding their usage across multiple products.

**Olivier Pomel**
*Co-Founder, CEO & Director*

Yes. I mean, look, if you look at these cohorts, and if you compare them in particular to the non-native cohorts we might have had like 5 or 6 years ago, like these are adopting many more products. So we see much broader product adoption and platform adoption with those customers than we saw 5, 6 years ago. So that's definitely something that is happening. Some of that is happening at renewal time, sometimes just happening organically as new needs emerge and as we spread to more teams inside those companies.

**David M. Obstler**
*Chief Financial Officer*

But I think it's important to clarify, these are renewals, but they are at higher levels than their previous commit, which is how they're getting the volume discounts. But we are renewing these customers at higher commits.

**Stewart Kirk Materne**
*Evercore ISI Institutional Equities, Research Division*

Okay. And David, I guess just for you, I mean, this would seem this cohort coming up for renewal earlier is all good over the long term. I'd imagine just visibility is a little bit harder from a sequential perspective. Is that sort of -- to Oli's comments on the long term, I'd imagine it's all good from a longer-term perspective, but it just creates a little more volatility for you on a sort of quarter out guidance perspective. Is that fair?

**David M. Obstler**
*Chief Financial Officer*

I would say it has a small amount of variation in the timing of the revenue path of these customers. The sequential here was really 3 things. It was really, one, that factor. Two, I think we mentioned the largest customer in the guide, which we're not predicting what they're going to do. And three, is we know about the small amount of seasonality. So it's one of the factors. We try to get a good sense of that. But of

course, as our RPO and as our contract, as we extend, it does create some volatility around the edges in the sequential nature.

**Olivier Pomel**
*Co-Founder, CEO & Director*

But again, it's all good news. Look, we are a usage-based model and when we deal with fast-growth customers, these customers themselves don't know exactly how they're going to grow over the next few months and what part of their business and how their consumption is going to change. So we happily work with them and make sure that we support their growth and we do everything that's right. The only time we think about it in difficult terms is when we prepare for these calls, we have to figure out, okay, so what part can we communicate, what part do we know, what part do we not know. But overall, these are all great news, great customers, great growth, a great adoption, amazing trends that are underpinning the industry as a whole. So we feel very good about it.

**David M. Obstler**
*Chief Financial Officer*

Yes. And this is similar to what we -- in the company's existence that you talked about. I mean our motion is to increase the commitment of customers, broaden the platform and extend the contract. So that's what we do all the time. It has some variability, but it's very similar to what we've been doing in the company throughout its history, and it's a good thing.

**Operator**
And our next question will come from the line of Peter Weed from Bernstein Research.

**Peter Weed**
*Bernstein Institutional Services LLC, Research Division*

You talked about it a little bit, but it's really exciting to hear about the release of the Bits AI, Scout and the strong early adoption. Could you help enumerate that a little bit? And maybe remind us how the pricing works there? Is it kind of directly monetized? Or does it really drive the adoption or use of Bits AI overall? It would be really exciting to kind of appreciate how broad this AI adoption is now and how it might move the needle on overall growth and perhaps become another major pillar for the business.

**Olivier Pomel**
*Co-Founder, CEO & Director*

Yes. So it's actually a bit of both. So look, it largely sits on the -- all of the data we acquire and the way we can process it and all the various products we have. So that's the core. But we also have a model where we charge for AI credits specifically for what Bits does on top of that. And the model we've chosen is a model where we stay fairly close to inference costs there so that we can pass on the savings as the models become more efficient to our customers, and so we can run many more investigations.

And the reason we've done that is the field is changing so fast, and we're improving the efficiency of those models so fast that it's not very productive to have a fixed price on say everything you do as you expect the volume of things you do to increase by orders of magnitude while the price goes down by orders of magnitude as well.

And so we are working on that model for that, which works really well. So long term, we expect both to monetize Bits directly as we do today and to drive broader adoption of the platform and deliver more value to our customers. In terms of what becomes the dominant monetization, it's hard to tell, and we don't really care. Again, we're fully usage-based. So as long as we get more usage, deliver more value, we'll be fine.

**Operator**
Our next question will come from the line of Koji Ikeda from Bank of America.

**Koji Ikeda**
*BofA Securities, Research Division*

I wanted to ask a question on Infrastructure Monitoring. So you guys disclosed it's now exceeding $2 billion in ARR and growing in the mid-30s. Could you quantify how different that growth rate is from when you last disclosed the scale of this business at $1.5 billion earlier this year. How much of that growth is coming from AI infrastructure? And how sustainable is that growth rate as GPU efficiency improves?

**Olivier Pomel**
*Co-Founder, CEO & Director*

So -- I mean, look, we've seen an acceleration of the business in general. And not just thanks to AI native, but also thanks to the rest of the customer base, so we have mentioned that, and that applies to Infrastructure Monitoring also. Today, that growth is not actually related directly to GPUs. GPUs are still largely on -- we see them for our customers on training. We have products there, but we don't monetize those products widely yet. And so for us, that's an opportunity we have in front of us. That's not something that's baked into the current growth and the current numbers.

**Operator**
Our next question will come from the line of Patrick Colville from Scotiabank.

**Patrick Edwin Ronald Colville**
*Scotiabank Global Banking and Markets, Research Division*

To me, one of the most interesting points in the prepared remarks is about Datadog training, your own frontier models. And then Olivier, your comments around this behavior is not mass market, but seeing how the market develops. I guess I want to ask you because you're at the forefront of how things are changing. How does this affect Datadog's use of proprietary tokens internally or for customers versus open weight tokens? And then I guess, David, if this shift that Olivier is describing picks up steam, what might that mean for Datadog profitability that it nicely increased this quarter, but is there a potential that this might help the bottom line as well?

**Olivier Pomel**
*Co-Founder, CEO & Director*

So look, what we see internally and what we see with customers is -- there's a mix of models that are being used. So I mean I can say we include both customers and us into that collective we because the behaviors are fairly similar. We are big consumers of the top labs, top models, but we also consume open weight models, a growing amount of open weight models, and we train our own, post-train our own.

And I think it's probably representative of where the industry as a whole is going. We do that because there are certain things you can only get from the top-of-the-line models that are being generally trained by others. And there's a number of things that are good enough and very cost efficient with open weight models, you can just pick up and run. And there's finally a number of other things that you can make way more efficient with your own data, your own workflow and your own post-training, which is what we've done with some of the models we announced.

So I expect that mix to remain about the same. In terms of the impact on margin, look, at the end of the day, you get some control on what you spend your money on. If it's very proprietary to you and you can deliver superior economy, you should absolutely do so. If not, like if it's very undifferentiated and applies broadly in some parts of your business that are the same to millions of other businesses, then it's fine probably to use off-the-shelf models and keep consuming those.

**Operator**
Our next question will come from the line of Alex Zukin from Wolfe Research.

**Aleksandr J. Zukin**

*Wolfe Research, LLC*

Oli, maybe a quick one for you and David, a follow-up. Just to what extent are you seeing the demand for Agent Observability from both AI natives and the core driving demand for offerings like your Traces product? And what is the -- like what has changed here longer term just given the kind of elevated adoption levels? And maybe also any comments on cybersecurity and the plans there?

**Olivier Pomel**
*Co-Founder, CEO & Director*

Yes. So we see growing demand for Agent Observability. And what -- so there was always interesting demand from the AI natives et cetera. But what's most interesting here is the non-AI natives are consuming more and more. And -- because they're getting deeper into AI, we see real demand, real volumes and really exciting -- really exciting pipeline for products actually. So we think that this is now crossing the chasm, so to speak, and that's exciting for this product, even though we didn't mention any hard numbers today on that.

On cybersecurity, look, there's so many new opportunities with AI. We're very happy to see that we can land the full platform with large customers. We mentioned that for a defense agency. We mentioned that also for a couple of other wins when we landed not just observability, but our security platform at the same time. So we feel good about the market fit there, and we feel good about the way the whole retooling and rebuild that is taking place with AI is playing in our favor.

**Aleksandr J. Zukin**
*Wolfe Research, LLC*

Perfect. And David, maybe for you, just a clarification. If we think about -- first of all, it's great to hear about the renewal with your largest customer. And I think we totally understand the guiding for the minimum for Q4. Any pearls of wisdom on how to think about modeling that for '27? Because presumably, that's a 1-year contract. And should -- is that something we should think about kind of maintaining at some level?

Is that something that you guys are going to take out of the model? And then if I think about how you're guiding with respect to some of those other early renewals for the AI natives that got renewed, are those also at the minimum commitment level? Are those slightly above? Just any help on those 2 items, I think, would be helpful for.

**David M. Obstler**
*Chief Financial Officer*

Yes, reminding everybody -- yes, yes, sorry -- yes, on the -- we'll give guidance for 2027 on our year-end call. I think we've said in the last couple of quarters that we're guiding at whatever the contractual minimums, and we've repeated that this quarter. We'll give some more flavor on that when we do '27.

And in terms of the rest of the business, what we've been -- what we do is, and we've said this over the many years, is we take the growth rate of those customers and we discount that. So that would be how we apply it across the customer base in taking the most recent organic growth and then discount it. And we would certainly be doing that for both the early renewal customers and our broader customer base.

**Olivier Pomel**
*Co-Founder, CEO & Director*

All right. And I think that's it for our questions today. So thank you all for listening. I also want to thank everybody at Datadog for a great Q3. I know we are right in the middle of an even greater Q4. We're all very, very busy. We're meeting -- I'm going to meet a lot of our customers at the largest AWS conference in a few weeks. And we hope to have many more exciting product announcement to share soon. So thank you very much all, and see you next quarter.

**Operator**
Thank you for your participation in today's conference. This does conclude the program. You may now disconnect.