

DDOG Earnings Call Transcript

Date: 2025-11-06

Quarter: 3

Operator: Good day, and thank you for standing by. Welcome to the Third Quarter 2025 Datadog Earnings Conference Call. [Operator Instructions] Please be advised that today's conference is being recorded. I would now like to hand the conference over to your first speaker today, Yuka Broderick, Senior Vice President of Investor Relations. Please go ahead.

Yuka Broderick: Thank you, Martin. Good morning, and thank you for joining us to review Datadog's third quarter 2025 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. During this call, we will make forward-looking statements, including statements related to our future financial performance, our outlook for the fourth quarter and the fiscal year 2025 and related notes and assumptions, our gross margins and operating margins, our product capabilities and our ability to capitalize on market opportunities. The words anticipate, believe, continue, estimate, expect, intend, will and similar expressions are intended to identify forward-looking statements or similar indications of future expectations. These statements reflect our views only 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, 2025. Additional information will be made available in our upcoming Form 10-Q for the fiscal quarter ended September 30, 2025, 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 the tables in our earnings release, which is available at investors.datadoghq.com. With that, I'd like to turn the call over to Olivier.

Olivier Pomel: Thanks, Yuka, and thank all of you for joining us this morning to go through our results for Q3. Let me begin with this quarter's business drivers. We have seen broad-based positive trends in the demand environment with an ongoing strength of cloud migration and digital transformation. Against this backdrop, we executed a very strong Q3 both in new logo bookings and usage growth of existing customers. As a notable inflection, we saw acceleration of year-over-year revenue growth across our non AI customers. And the sequential usage growth for non-AI existing customers was the highest we have seen going back 12 quarters. This growth was broad-based as our customers are adopting more products and getting more value from the Datadog platform. We also experienced strong revenue growth for our AI native customers and a broadening contribution to growth among those customers. There, too, we saw an acceleration of growth in our AI cohort in Q3 when excluding our largest customer. Looking at new business. Contribution from new customers increased in Q3 in both the amount of new customer bookings as well as the revenue contribution from new customers. And as usual, churn has remained low with gross revenue retention stable in the mid- to high 90s, highlighting the mission-critical nature of our platform for our customers. Regarding our Q3 financial performance and key metrics. Revenue was \$886 million, an increase of 28% year-over-year and above the high end of our guidance range. We ended Q3 with about 32,000 customers up from about 29,200 a year ago. We also ended with about 4,060 customers with an ARR of \$100,000 more, up from about 3,490 a year ago. These customers generated about 89% of our ARR. And we generated free cash flow of \$214 million with a free cash flow margin of 24%. Turning to customer adoption. Our

platform strategy continues to resonate in the market. At the end of Q3, 84% of customers were using 2 or more products, up from 83% a year ago. 54% of customers were using 4 or more products, up from 49% a year ago. 31% of our customers were using 6 or more products, up from 26% a year ago and 16% of our customers were using 8 or products, up from 12% a year ago. Digital experience is an example of an area within our platform where our rapid piece of innovation is turning into tangible value for our customers. Our digital experience products include RUM or Real User Monitoring, to observe and improve application behavior in mobile and web apps, synthetics to stimulate user flows and proactively detect user facing issues and product analytics to help users connect application behavior to business impact. Over the years, we built our product breadth and depth in this area, and that is being recognized in the marketplace. For the second year in a row, Datadog has been named the leader in the 2025 Gartner Magic Quadrant for Digital Experience Monitoring. We are pleased that today, these digital experience products together exceed \$300 million in ARR. And this includes, in particular, a very fast ramp for product analytics, which has already seen adoption by more than 1,000 customers. We also want to call out our security suite of products where we are executing and accelerating growth. Security ARR growth was in the mid-50s as a percentage year-over-year in Q3, up from the mid-40s we mentioned last quarter. We're starting to see success in including Cloud SIEM in larger deals, and we'll get back to that in a bit in our customer examples. And we're seeing positive trends beyond Cloud SIEM, including fast uptake of good security and an increasing number of wins in cloud security. Overall, we saw year-over-year growth acceleration in each one of our security products. Moving on to R&D.; We continue to deliver on what is a very ambitious AI road map. We are seeing high customer interest in our Bits AI agents, which we announced at our DASH user conference in June. We have now onboarded thousands of customers for preview access, the Bits AI SRE agent. And as we prepare for general availability, we are getting very enthusiastic feedback on the time and cost savings enabled by Bits AI. As RUM user recently told us, with Bits AI SRE being on call 24/7 for us, meantime to resolution for our services has improved significantly. For most cases, the investigation is already taken care of well before our engineers sit down and open their laptops to assess the issue. And this is not an isolated comment. We see the potential here for our agents who radically transform observability and operations. In LLM observability, we recently launched LLM experiments and playgrounds for general availability, helping teams to rapidly iterate on LLM applications and AI agents. We also launched custom LLM as a judge evaluations for general availability, which lets customers write evaluation prompts to assess application quality and safety. As an illustration of growth and adoption in the past few months, the number of LLM spans customers are sending to Datadog has more than quadrupled. And we are seeing a lot of interest in the Datadog MCP servers. Our MCP server acts as a bridge between Datadog and AI agents, such as Codex OpenAI, Claude by Anthropic, Cursor, GitHub Copilot, Goose by Block and many more. Our preview customers are using real-time production data context to drive trouble shooting, root causes analysis and automation in these agents. One user told us, the Datadog MCP server is a great tool. It enables me to get the last 5 of my app and follow the spans and traces all the way to the root cause. I have never been more hooked on Datadog. So we see MCP adoption as a great way to cement Datadog even further into our customers' workflows. Finally, we continue to see rising customer interest for next-gen AI observability with over 5,000 customers sending us AI data to one or more of AI integrations. On the topic of integrations, we are very proud to now support over 1,000 integrations, which we believe is unparalleled in our space. By using our integrations, customer call it otherwise disparate data sources across Datadog products for deeper analysis. We can see from a customers usage that this is a critical part of the Datadog platform. Our 32,000 customers use more than 50 integrations on average, while customers spending over \$1 million annually with us use more than 150. And most importantly, as tech stack evolves, we continue to update and expand our integrations. So our customers can use Datadog to deploy new technologies with confidence. Last but not least, I wanted to give a shout out to our AI research team for the amazing work they have published. Our TOTO OpenWave time-series forecasting model has been one of the top downloads on Hugging Face over the past few months, and that is across all categories. It is very impactful as, among other things, the high quality of this work allows us to attract world-class AI researchers and engineers. Now let's move on to sales and marketing. We had a number of great new logo wins in customer expansion this quarter. So I'll go through a few of them.

First, we landed a 7-figure annualized deal with a leading European telco, our largest ever land deal in Europe. This company's previous setup was expensive, inefficient and wasn't scaling to meet their needs. By using Datadog, they expect to save over \$1 million annually on tool cost alone, along with millions of dollars more in reduced operation costs, lower engineering time and avoidance of revenue loss. They will adopt 11 Datadog products to start, and we consolidate more than 10 commercial and open source tools. Next, we landed a 7-figure annualized deal with a leading financial risk and analytics company. The company's fragmented tooling has led to major incidents that sometimes took multiple days and hundreds of engineers to resolve. They plan to start with 11 Datadog products including On-Call, Cloud SIEM and Bits AI, and will replace 14 commercial open source and hyperscale observability tools. Next, we landed a 7-figure annualized deal with a Fortune 500 technology hardware company. This is an exciting win for new -- sorry. This is an exciting win for our new go-to-market motions, targeting the largest and most sophisticated companies in the world. Datadog has been chosen as their strategic observability partner, and we are displacing commercial tools across availability, cloud team and incident response. This customer is starting with 14 Datadog products. Next, we signed a 7-figure annualized expansion with a Fortune 500 financial services company. This customer has pockets of siloed teams and data, including one business unit, which manually hosted and maintained 93 separate instances of open source tooling. With this expansion, this company will adopt 15 Datadog products, including all 3 pillars in all of their business units. They will also replace their SIEM solution with Datadog Cloud SIEM in a 7-figure land deal for Cloud SIEM. And by bringing all their telemetry data into the Datadog platform, they expect better insights for their adoption of Bits AI SRE Agents today and Bits AI [indiscernible]. Next, we signed a 7-figure annualized expansion with a Fortune 500 heavy equipment company. With this expansion, this customer will replace its open source log solution with Datadog log management and Flex logs. They plan to adopt LLM Observability and their IT team is using cloud cost management to improve cost visibility and governance. Next, we will come back a leading vertical SaaS company with a 7-figure analyze deal. By returning to Datadog, this customer benefits from our alignment with open telemetry and we'll implement the incident and reliability processes that they were unable to execute on previously. Next, we signed a 7-figure annualized expansion with a major American carmaker. This customer is adopting Datadog products faster than previously expected and this agreement supports the higher usage. With this expansion, they will adopt Datadog incident management and On-Call solution company-wide for a total of 5,000 users who support operational continuity across the business. Finally, we signed a 9-figure annualized expansion with a leading AI company. This company has been a long-time Datadog customer and has expanded their usage of multiple products, securing better economics for a higher commitment with an early renewal. Speaking of AI customers, we continue to help AI native customers big and small to grow and scale their businesses. And we continue to see this group broaden in number and size with more than 500 AI native companies in this group, but 100 of which are spending more than \$100,000 annually with Datadog and more than 15 who are spending more than \$1 million annually with us. While we know there's a lot of attention on this cohort, we primarily see it as an indication of what's to come as companies of every size and every single industry incorporate AI into their cloud applications. And that's it for another very strong quarter from our go-to-market teams, who are now very hard at work as we have a really exciting pipeline for Q4. Before I turn it over to David for a financial review, I want to say a few words on our longer-term outlook. There is no change to our overall view that digital transformation and cloud migration are long-term secular growth drivers of our business. Meanwhile, we are advancing rapidly in AI, where we are incredibly excited about our opportunities. We're building a comprehensive set of AI Observability products to help our customers tackle the higher complexity that comes with these technologies. And we are building AI into Datadog, and I spoke earlier about the excitement our customers have for our Bits AI agents. The market opportunity in cloud and AI is expected to grow rapidly into the trillions of dollars and companies of every size and industry are looking to adopt AI to deliver value to their customers and drive positive business outcome. So we're moving fast to help our customers develop, deploy and grow into the cloud and into the AI world. With that, I will turn it over to our CFO. David?

David Obstler: Thanks, Olivier. To start, our Q3 revenue was \$886 million, up 28% year-over-year and up 7% quarter-over-quarter. To dive into some of the drivers of our Q3 revenue growth. First, overall,

we saw sequential usage growth from existing customers in Q3 that was higher than our expectations and the strongest in 12 quarters in our non-AI native customer base. We saw year-over-year growth acceleration broadly across our business, including in new logos and existing customers, both enterprise and SMB, with customers across our spending bands, big and small, and customers in a wide variety of industries. Next, we saw strong and accelerating contribution from new customers. New logo annualized bookings more than doubled year-over-year and set a new record driven by an increase in average new logo land size, particularly in enterprise. We believe we are starting to see the benefits of our growth of sales capacity. And we are seeing new logos ramping faster, contributing more to revenue growth. The portion of our year-over-year revenue growth that related to new customers was about 25% in Q3, up from 20% in Q2. Next, our AI native customers continue to exhibit rapid growth, while more customers in this group are growing to be sizable customers. As Olivier discussed, we extended the contract of our largest AI native customer. In addition, we now have more larger AI customers, including 15 of them spending \$1 million or more annually with Datadog, and about 100 spending more than \$100,000 annually. Year-over-year revenue growth from our AI native customers, excluding the largest customer, again, accelerated in Q3. In Q3, this group represented 12% of our revenue, up from 11% last quarter and about 6% in the year ago quarter. I will note that over time, we think this metric will become less relevant as AI usage in production broadens beyond this group of customers. Our year-over-year revenue growth also accelerated amongst our non-AI native customers. In Q3, our revenue growth, excluding the AI native customer group, was 20% year-over-year, accelerating from 18% year-over-year in Q2, and we have seen this trend of accelerating growth continue in October. Regarding retention metrics. Our trailing 12-month net revenue retention percentage was 120% similar to last quarter and our trailing 12-month gross revenue retention percentage remain in the mid- to high 90s. And now moving on to our financial results. Our billings were \$893 million, up 30% year-over-year. Our remaining performance obligations or RPO was \$2.79 billion, up 53% year-over-year, and current RPO growth was in the low 50s percentage year-over-year. Our strong bookings contributed to this acceleration of RPO. We continue to believe that revenue is a better indication of our trends in our business than billings and RPO. And now 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. First, gross profit in the quarter was \$719 million, and our gross margin was 81.2%. This compares to a gross margin of 80.9% last quarter and 81.1% in the year ago quarter. As previously mentioned, we continue to see the impact of our engineers cost-saving efforts in Q3 as they deliver on our cloud efficiency project. Our Q3 OpEx grew 30% -- 32% excuse me, year-over-year, down from 36% last quarter. We continue to grow our investments to pursue our long-term growth opportunities, and this OpEx growth is an indication of our execution on our hiring plan. Q3 operating income was \$207 million for a 23% operating margin compared to 20% last quarter and 25% in the year ago quarter. And now turning to our balance sheet and cash flow statements. We ended the quarter with \$4.1 billion in cash, cash equivalents and marketable securities and cash flow from operations was \$251 billion in the quarter. After taking into consideration capital expenditures and capitalized software, free cash flow was \$214 million for a free cash flow margin of 24%. And now for our outlook for the fourth quarter and the fiscal year 2025. First, our guidance velocity overall remains unchanged. As a reminder, we based our guidance on trends observed in recent months and imply conservatism on these growth trends. So for the fourth quarter, we expect revenue to be in the range of \$912 million to \$916 million, which represents a 24% year-over-year growth. Non-GAAP operating income is expected to be in the range of \$216 million to \$220 million, which implies an operating margin of 24%. Non-GAAP net income per share is expected to be in the range of \$0.54 to \$0.56 per share based on approximately 367 million weighted average diluted shares outstanding. And for the full year -- fiscal year 2025, we expect revenues to be in the range of \$3.386 billion to \$3.390 billion, which represents 26% year-over-year growth. Non-GAAP operating income is expected to be in the range of \$754 million to \$758 million, which implies an operating margin of 22%. And non-GAAP net income per share is expected to be in the range of \$2 to \$2.02 per share, based on 364 million weighted average diluted shares. And finally, some additional notes on our guidance. We expect net interest and other income for the fiscal year 2025 to be approximately \$170 million. We continue to expect cash taxes in 2025 to be about \$10

million to \$20 million and we continue to apply a 21% non-GAAP tax rate for 2025 and going forward. And finally, we expect capital expenditures and capitalized software together to be 4% of revenues in fiscal year 2025. To summarize, we are pleased with our execution in Q3. We are well positioned to help our existing and prospective customers with their cloud migration and digital transformation journeys, including their adoption of AI. And I want to thank Datadog's worldwide for their efforts. And with that, we'll open the call for questions. Operator, let's begin the Q&A.;

Operator: [Operator Instructions] Our first comes from the line of Kash Rangan of Goldman Sachs.

Kasthuri Rangan: Appreciate it. Congratulations on the spectacular results and showing sequential improvement across the board. Olivier, I had a question for you. We've talked about GPU monetization versus CPU monetization. So how closer are we to the point where you can confidently expand and get your share of the customer wallet when it comes to whether it's training workload, inferencing workload on the GPU clusters, which are becoming more prevalent and increasingly a larger part of the compute build-out in the future? That's it for me.

Olivier Pomel: Yes. So we have products that are getting into the market now for GPU monitoring. But these don't generate any significant revenue yet. So all the revenues we've shared, like the acceleration, et cetera, that's not related to us capitalizing more on GPUs, that's a future opportunity.

Operator: Our next question comes from the line of Sanjit Singh of Morgan Stanley.

Sanjit Singh: Congrats on the acceleration in growth this quarter. Olivier, I wanted to talk about some of those enterprise trends you're seeing in sort of your non-AI cohort. What do you sort of put the improved performance in growth this quarter on? You mentioned that the sales productivity or the benefit from some of the sales investments starting to come online. Is there sort of an uplift in sort of the cloud migration trends as you're starting to see enterprise build more AI applications. I'd just love to get your perspective on the underlying trends in the enterprise and the mid-market business.

Olivier Pomel: Yes. I'd say there's 3 parts to it. One part is the demand environment is not -- is positive in general. I don't know that we see massive acceleration of cloud migration, but at least the environment is not pushing the other way. We know which happens from time to time. So that's point number one. Point number two is we've been growing sales capacity quite a bit, and we've created new go-to-market motions to go after the kind of customers who were not getting before. Like we've done quite a bit of investment over the past couple of years and we see that starting to pay off. As I said also, we feel good about Q4 in terms of pipeline on the sales side. So it's too early to tell yet. We still have to close those deals, but we feel good about the scaling of our go-to-market. And point number three is we have a number of products that we've been developing over the years. Some of them are early, some of them a little bit further along that are really clicking. We see -- we have a lot of success with getting large enterprises to adopt Flex Logs, for example. We have a lot of success in some of new products such as analytics that we mentioned on the call. We're seeing some large land deals with our cloud team. So all of that is contributing to the picture you're seeing today.

Sanjit Singh: And just as a follow-up on the AI observability opportunity. When you look at some of the independent software vendors that are releasing Agentic solutions, Agentic portfolios. A number of them are including observability as part of their sort of value proposition. Is there any work you think Datadog has to do to sort of infiltrate that market or make sure that customers look to Datadog as that Agentic monitoring capability as some of these independent software vendors try to bundle in observability into their solutions. I would love to get your perspective on that?

Olivier Pomel: Yes. I mean there's absolutely no doubt to us that the customers will even want a unified platform for observability for all of this. There's 2 parts to that. One is, historically, every single piece of software we integrate with, whether that's SaaS or things that customers on themselves, also has its own management control and observability control. But you're not going to log into [70] or in the case of customers we mentioned that they use 60 integrations for the smaller customers, 150 integrations for the larger ones. It's not practical to actually go and manage that separately. So we think all of that belongs in a central place, and that's the historical trend we've seen. We also think that you can't separate the AI parts from the non-AI parts of the business. So you're not going to look at your agents separately that you do at your web hosting and your database and your -- everything else you have in your stack. So all of that in the end will be attached to observability.

Operator: Our next question comes from the line of Raimo Lenschow of Barclays.

Raimo Lenschow: Perfect. Congrats from me as well. That sounded like an amazing quarter and nice to see it coming together. On the AI side, and I don't want to talk about the customer, but more the other ones, like 15 customers over 1 million. That's like a big number and 100 over 100,000. How do we have to think about the nature of those? Is this kind of -- are those kind of especially the bigger ones of those kind of model builders, but then even 15 is a big number. And over 100 sounds like this whole new application world that we've all been kind of waiting for starting to come together. Is that kind of what's going on there? Because it does sound quite exciting and much more broader than we thought.

Olivier Pomel: It's actually fairly broad. So there is model vendors, there's models -- model that can be the lens model that can be video, it can be sound generation, it can be all of the various parts of the stack you see as independent companies. It can be -- there's quite a few companies that do that work on the coding side. So coding assistants and vibe coders and everything in that range. Some of these are very new companies. Some of these are not very new companies, some of these started 5, 7, 8 years ago. And we're sort of not necessarily AI native from day 1, but very quickly, that would give them the growth they see today with the people to AI. So we see a little bit of that. We have companies that are other parts of the stack in AI on the, say, the [steady] side, the other components of the infrastructure. And we have other companies that are purely applications filled with AI. So we have a bit of everything in there. Like it's actually fairly representative of the space.

Operator: Our next question comes from the line of Mark Murphy of JPMorgan.

Mark Murphy: You had mentioned the expansion of the contract with your largest AI native customer and I believe you said better economics for a higher commitment. Can you speak to that because I would assume a higher commitment would carry a volume-based discount. I'm just trying to understand if. For some reason, if that was not the case here, what did you mean by better economics? And then I have a quick follow-up.

Olivier Pomel: Yes. I mean, this is without getting to the detail of any specific customer like this is the motion is always the same, like customers grow, they commit to more, they get better prices. So you see, like a, again, talking about customers in general, you see both of usage, drops in revenue as customers renew and get higher commit and a better price and then usually growth after that for those customers. That's the motion that we've had. We have about 30,000 customers so far.

Mark Murphy: Okay. And the -- what in that, so the better economics part of it is just where it's going to be netting out like 12 months down the road? Is that what you mean?

Olivier Pomel: Well, the bigger economics means you commit tomorrow, you get a better price. And as we -- remember, we have a usage model. So we charge people every month on what they use at the price we agreed. So if you get better economics, and your usage is somewhat similar month to month, one month and the next that you pay less, but the overall backdrop of our business is increased consumption.

Mark Murphy: Okay. And then as a quick follow-up, Olivier, the acceleration that you saw in the security growth is pretty noticeable too. We recall, I think about 6 months ago, you had ramped up and engaged a lot more of a channel partners, which is a key ingredient to grow in the security business. Is it a function of that? Or is there a mindset change happening out there where customers want observability to be the central point of collection so that all the security teams and the ops teams are working with the same set of metrics and logs and tracers?

Olivier Pomel: Look, I think it's a number of things. Definitely, we've been investing in the channel, and that's certainly helpful to do the security business as a whole. The win -- the big win we mentioned on security that we mentioned a couple of wins in Cloud SIEM. These tends to be more related to product maturity. The strength of our underlying platform, especially when it comes to technology like Flex Logs, for example. And the fact also that we've been learning how to properly go to market for security. And I think we see things clicking in a way that is exciting.

Operator: Our next question comes from the line of Fatima Boolani of Citi.

Fatima Boolani: Oli, I'll start with you and I have a follow-up for Dave. On the On-Call product, Oli, how do agentic advancements in general detract or enhance the value proposition here? And I'm very simplistically thinking about the core nature and value proposition of the On-Call product intelligently routing, requests for remediation, right? So how do you just broader advancements in AI, help beef up

and/or detract your ability to monetize this product? And then just a follow-up for David, please.

Olivier Pomel: Well, I mean, if you zoom out, we entered the field with On-Call because we wanted to own the end-to-end incident resolution. So we wanted because we before that, we were detecting the incidents and sending the alerts, and then we were pretty much where the resolution happened after that. Customers were spending their time in data to diagnose and understand what was going on. So we wanted to own the full cycle. . And we thought that with AI, in particular, we'd have the ability to do things if we are on the whole cycle that we couldn't do otherwise. So what you see right now is, I mean, this resonates with customers, they adopting to product. We've mentioned like some exciting customers with say, [one] with 5,000 seats for On-Call, which is very exciting. But in the future, there's many more things we can do in working on for that product. If we both detect incident and notify, we can do some sort of things such as even predicting the incident and notifying early or rerouting early or telling people before the incident actually takes place, how they can potentially fix it. So these are all things we're working on. I mean, look, if you look at the various product announcements we've made, whether that's Bits AI or SRE or the time series forecasting model we have released. When you assemble all that, you get to a very, very interesting picture of what we can do in the future. So we're excited by that. Our customers are excited by the vision there too, and that's why these products are successful.

Fatima Boolani: Appreciate that. David, on net retention rates, why aren't we necessarily seeing more upward pressure on the metric, just given the strength of expansionary bookings that you alluded to in the quarter from the installed base. And I mean I suspect it's because it's a trailing 12-month metric. But any directional color you can just share on that. And any high-level commentary on some of the non-AI native net retention rate trend behavior?

David Obstler: Yes, you've nailed it. It's a trailing 12 months, it's a number that's rounded, it has the dynamics that you might expect in that the growth of the non-AI natives has been, as we mentioned, a combination of landing and expanding at higher rates than we've seen in recent quarters. So if that continues as you go into a trailing 12-month metric, you see a directional movement.

Operator: Our next question comes from the line of Eric Heath of KeyBanc.

Eric Heath: Oli, David, Bits AI seems like a really exciting thing out of Dash. And I know it's still in preview, but you mentioned there's a lot of interest there. So I'm just curious how you think about the agentic opportunity with Bits AI. How meaningful this can be for 2026 as a differentiator versus competition and also as a revenue contributor?

Olivier Pomel: Yes. So -- I mean, look, it's super exciting. The feedback is very good on it. I mean, we've been collecting all the -- so I read one quote, we have dozens that look just like that, that was sent to us by customers. And so that's very, very exciting. We also started having some customers buy and come to it to just to show value and to make sure we're on to the right product mix. And so we feel good that this is something that is high quality and we can monetize. In terms of the impact for next year, on the packaging side, I'm not completely sure yet whether the biggest impact will be seen from what we charge Bits AI itself or for the rest of the platform, that it gets benefits from the differentiation of Bits AI. I think that's more of a broader question of packaging and monetization of AI. And remember that we have a product that is usage based. So anything that drives usage up and adoption from customers is good for us and is very, very monetizable. But what we can tell is this is differentiating, this is good. It works significantly better than anything else we've seen or heard of in the market, and we are doubling down on it. We have many, many teams now working on deepening Bits AI SREs to making sure it goes further into the resolution doesn't just point to the issue, but fixes the code that all these kind of things working hard on that. We're also working on breadth, making sure that we train it on many more types of data, many types of sources, sometimes even systems that are observe the systems that are not dialed up, so we can cut across to other systems our customers are using. So we are very, very aggressively developing Bits AI SRE. It's resonating very well in the market.

Operator: Our next question comes from the line of Gray Powell of BTIG.

Gray Powell: Congratulations on the great results. So maybe just like taking a step back, if we go back to the beginning of the year, Datadog was expecting 19% revenue growth. It looks like you're tracking to something over 26% growth now, and that's just the high end of your guidance. So I guess my question is, what surprised you the most this year? And then just how do you feel about the sustainability of those drivers as you look forward?

Olivier Pomel: I mean, look, the -- so first, I apologize for over delivering on the results. We might do it again, but we'll see. I think the biggest surprise for us has been that -- so AI in general has or AI adoption has grown faster than we thought it would at the beginning of the year. So we've seen that across our AI cohort. We've seen also that we got some of our new products and new, like the changes we're making on the go-to-market side to click perhaps earlier than we would have thought otherwise. So all in all, we saw the leading part of the business with AI growth faster, not the lagging but the slower growing, more traditional part of the business also accelerate and that gets us where we are today.

David Obstler: And I'd add, we have a good demand environment and we've been investing whether it be in the products that Oli's been talking about or in the sales capacity we made clear that we were in investment and we're seeing those investments pay off.

Operator: Our next question comes from the line of Koji Ikeda of Bank of America Securities.

Koji Ikeda: Just one from me here. I wanted to ask a question on the inflection in the non-AI native growth and how to think about the areas of strength in this cohort. Is it coming from your largest enterprises? Is it coming from a certain type of customer? Is there a common theme in the workloads that you're seeing or the products that are being added on that is driving that strength? Or is it just really just broad-based? What I'm trying to get out here is I'm really trying to understand more the durability of this growth reflection.

Olivier Pomel: So it is broad-based. And I think, again, speaks to a couple of things. It speaks to the fact that, in general, the demand environment is good. Though I would say, there's been a very, very high growth of hyperscaler revenue like over the past an acceleration for the hyper scalers in general. A lot of that is GPU related, but the growth we're seeing here and the exception we're seeing here is largely not GPU-related. It's a little bit of it, but not a ton of it. So that's not exactly what you've seen with some of the other vendors there. One reason this is broad-based is these are the same products we sell to all customers, and this is largely the same go-to-market organization that we have a few segments, but -- and we've been doing well executing there. I think we've invested quite a bit in product, and we keep and we will keep doing it, and we see the results of that.

David Obstler: Yes, I want to -- I'll add that it's across the customer base, enterprise SMB. And when we look at it, it's not just an AI SMB. If you remove the AI companies, you still see a strengthening SMB demand cycle going on. And unlike in previous periods, it also is across spending ranges. We're not seeing larger spenders or smaller spenders. We're just seeing a broad trend of improved demand across the spending trends.

Olivier Pomel: And remember that for us, SMB is, any company of less than 1,000 employees. It includes a lot of very legitimate and growing businesses. It's not...

Operator: Our next question comes from the line of Ittai Kidron of Oppenheimer & Co.

Ittai Kidron: Congrats guys. Really great numbers. Oli, in your answer to one of the questions and kind of going into the drivers behind the upside. You've talked about sales capacity increase. You didn't talk much about sales efficiency. Is there a way you can give us some color on where do you stand on percent of salespeople that are hitting quota, where does that ratio stand relative to historical patterns for you guys? And as you approach '26 year, do you anticipate any material changes in the comp structure just given the breadth of product and the list of opportunities, how do you get people focused?

Olivier Pomel: Yes. So we feel good about the sales productivity in general. And the rule generally, you grow by scaling capacity and maintaining productivity, it's hard to drive both up at the same time. And remember, if you want to go to 10x, you can do that by scaling if you can't really do it by improving productivity, so you have to scale. And we've been doing that, and we've been successful at it so far. In terms of the complaints that, look, we keep changing the way we compensate and the way we manage the sales force in general to make sure we have the right focus. One of the gifts of a business like ours is that we see -- we have a very heavy land-and-expand model. And so we get a lot of growth from existing customers. The challenge it creates on the other hand is how do we get to focus the sales force on the newer customers, the smaller ones and the new ones because it is more work to get an extra dollar for a smaller customer or for a new one and it is from an existing one that they already have scale. And so a lot of the tweaks we met to our comp plans relate to that. Who do we make sure we direct our attention and we reward people for what is going to generate the most long-term growth for

us. And we've made a number of changes. I won't go through them, these are our internal changes. But we had a number of changes this year, we see a number of them pay off. Another thing I mentioned on the call was you mentioned a win for one of our new go-to-market motions and that specifically getting in place multiyear plans to go after some larger customers that are tougher to land than what we've done in the past. And sometimes, it takes more than a year to land certain types of customers. And the problem is if you comp plan only has a 1-year horizon, like it doesn't give a great incentive for the sales force to go after those customers. And so we cordoned off a few of those companies who have special plans to go after that, and we're starting to see success with that too. It's just an example.

Operator: Our next question comes from the line of Andrew Sherman of TD Cowen.

Andrew Sherman: Great. Congrats. I know you have a team focused on the Fortune 500, where there's still a lot of white space for you. Curious to hear how the team is ramping to productivity. Did that help drive some of the strong new logo bookings and can this contribute even more next year?

Olivier Pomel: Yes. I mean, look, the team is not new, right? I mean, we've been focusing on that for many years, and we're tracking well. One thing I was mentioning just before was one challenge even in the Fortune 500 is to make sure that we focus on landing new customers and make sure that there's the right amount of sales attention and reward for the landing a customer even if it's for a small amount, and I think we've done well. I mean again, we can comment on that again after the next quarter when we have a full year of our new clients that have been validated. But so far, we feel very good about it.

Operator: Our next question comes from the line of Alex Zukin of Wolfe Research.

Aleksandr Zukin: Congrats on dropping some truly inspiring quotes in the script. Maybe Oli, one for you and then I have a quick follow-up for David. Just the duration of this acceleration of the non-AI cohort. It seems like from all your forward-looking metrics, whether it's billings, RPO, CRPO. Those were, again, really, really strong how long do you think we should think about the duration of this trend of this non-AI acceleration?

Olivier Pomel: Well, we're a consumption business. So we -- the hardest thing to understand is what the future is going to look like for consumption. The way I would say it is we feel very good about it at the midterm, long term. Now with ebbs and flow in any given month or quarter, that's harder to tell. And again, that's what we've seen through the life of the company. So we feel very confident about the motion in general for digital transformation and cloud migration is steady. And sometimes it slows down a little bit, but it reaccelerates after that. And we see that key going on for a very long time.

Aleksandr Zukin: Okay. And then maybe, David, for you, look, gross profit dollar acceleration while you're seeing your largest customer kind of get better unit economics is also inspiring to see how should we think about the progression of gross margins and gross profit dollar growth, particularly as you continue to also see the AI cohort acceleration.

David Obstler: Yes, there's a couple of things. I think we've mentioned that we've been focused and have focused over the many years on the efficiency of our cloud platform. We have significant engineering efforts around cost of sales and delivery of value. And so we've been able to deliver on that. We also have a very broad customer base distributed in terms of volume. So as customers get larger and maybe get volume discounts, we have a number -- a lot of customers coming in, it's smaller, so that balance there. And then in terms of the sort of the future -- I'll repeat what we've always said that we've been running the company with a gross margin plus or minus 80%, we've given that range and not changed it, and we watch it. And it gives us mixed signals in terms of efficiency, how we're operating, it gives us good signals in pricing and things like that and I wouldn't change the comments we made over the many years about looking at that and then developing operations and strategies around that.

Operator: Our next question comes from the line of Ryan MacWilliams of Wells Fargo.

Ryan MacWilliams: Just one for me. On the large AI contract expansion that you provided commentary on, is there any way we can think about the contribution change from this customer over the next few quarters?

David Obstler: No. I mean we don't provide that kind of information on individual customers. We're trying to give a picture of the overall business. Generally, I think as Oli mentioned, on our larger customers, we have a motion of the expansion of volume and then we talk when we work on the term and the volume-based pricing, but we don't give guidance like that on individual customers.

Operator: Our next question comes from the line of Mike Cikos of Needham.

Michael Cikos: I just wanted to come back to it, Oli, for the non-AI native strength, I know we've kind of hit on this a number of times, whether it's road map sales capacity execution, but like kudos on the numbers here? I'm just trying to get a better sense of the why now. Is it just a composite of all those different pieces clicking together this quarter? Or is there anything more to unpack there? And then I have a follow-up for David.

Olivier Pomel: Again, I don't think there's a lot more to unpack there. And I know it's boring in a way, but it's also the way we've been growing for the past 15 years, really. So that's a -- I would call it the usual.

Michael Cikos: Awesome. Awesome to hear. Okay. And then for the follow-up to David. David, I don't want to take anything away from the Q3 results you guys just posted, and we obviously have the strong guide here for Q4. But I just can imagine myself a month from now starting to get inbound from certain folks asking about the holiday season and the fact that we have the holidays landing on weekdays in Q4 here. Can you just kind of discuss how you thought about constructing guidance for this Q4 year?

David Obstler: Yes. We have years of experience of analyzing the day-by-day patterns. In the holidays, we know that the holiday period ends up in the usage side because of vacation holidays, and we incorporate that into our guidance. We, I think, evolved a lot over the years and sort of days adjusted types of days, et cetera. And so we would be incorporating that like we've incorporated in other years. If there are differences in this calendar period, we incorporate that as always.

Operator: Our next question comes from the line of Karl Keirstead of UBS.

Karl Keirstead: Okay. Great. I'll ask one for David and one for Olivier. David, first of all, congratulations on the extension of the larger contract, I think everybody on the line is applauding that. I know you're reticent to get into any details, but maybe I could try. Are you able to clarify whether that was a 1-year deal or multiyear? And then related to that, David, what is the contribution to CRPO from that deal, which I presume landed in your CRPO number. If it is a 1-year deal does the entirety of that contract contribute to the sequential CRPO performance in the quarter? So that's it for you, David. And then Olivier, maybe I'll just ask both at once. Some of the very large AI natives are beginning to diversify to utilizing Oracle's OCI and Stargate. And I'm wondering what's the opportunity for Datadog to essentially follow that behavior and begin scaling on Oracle's target or because a lot of what Oracle is doing with the AI native is training clusters, perhaps that near-term opportunity is more limited.

David Obstler: Yes. On the first point, I think we give a lot of examples and our motion, which our customers would be following, including that one would be -- we fix out annual plus commits. We're not commenting on individual contracts here, but it would follow a typical path to other types of contracts. So that's what we would do.

Olivier Pomel: Yes. And on the OCI, look, this is -- we've built an OCI integration, and then we see more demand from customers on OCI. Some of the things we see like the targets, et cetera, like these are extremely custom build out, like I don't know -- they're not necessarily exactly cloud because they are custom built for specific customers. So the opportunity there is more remote today. But it's -- again, one company is that it's a not fantastic opportunity to product type. But if 10, 15, 20, 50 companies start using that, then that really becomes a commercial opportunity. And so we're very much plugged into all of that. And we go basically where our customers are.

David Obstler: I think you mentioned about the RPO. I think in this case, we've mentioned this current and the total is roughly the same, and there wouldn't be anything in that contract that would have been materially around of those numbers. Those numbers, I think we mentioned are produced from the bookings growth more generally and not from that particular contract.

Operator: Our next question comes from the line of Jake Roberge of William Blair.

Jacob Roberge: Yes. Just on the recent go-to-market investments, obviously, it seems like there's been a lot of traction thus far with those. So I'm curious if there are any areas like security or the new logos or upmarket that you could look to lean even deeper into just given the growth that you've seen here.

Olivier Pomel: Yes, definitely. And there are some things we didn't do this year that we'll definitely go to the next year. So there's a number of things we are -- we're in Q4, right? So we're in the middle of planning for next year, and we basically will keep scaling what's working, stop doing some of the things that are not conclusive and then try to do more things. That's the way it works. Interestingly enough,

building a go-to-market is not that different from building software like you experiment together data you see what's working was not working and you build the systems.

Jacob Roberge: That's helpful. And then just on the new Bits AI Agents, can you just talk about the early feedback that you've gotten for those solutions and maybe how the engagement with those agents as compared to kind of the ramp of security Flex Logs. I know, obviously, much earlier days, but just how it compared when those were still largely in the preview phase?

Olivier Pomel: I mean look, the Bits AI Agent is -- it really has a growth factor for customers. So what works really well is and we've seen that number of times, like we set it up for them. It's running on their alert and they go through an outage and they still go to the motion, so they still go -- they still set up a bridge and they have 20 people and they spend 2 hours and in the end, they have an idea what went wrong. And then they go to Datadog and they see, oh, there's an investigation that had run. And 3 minutes into the outage, it got the same conclusion that we got 2 hours later with 20 people on the call. And that's completely eye-opening for customers when they see us. And we have -- so that's why we get many quotes about it. So now there's more we need to do there, like customers say, "Oh, it's great. Now can it make this fix for me? Can you do this? Can you do that? Can you support that other system that right now, you can't actually set it up for. So we have a very, very full road map of things we need to do, and we're doubling down on it. We also shipped -- I mean this one is in preview, but we shipped a security agent that looks at vulnerabilities and looks at security signals and those triad that basically look at trying to investigate what might be benign or what might be a real issue. We also are getting very, very positive feedback for that. And in fact, that's what helped us win some large land deals for our Cloud SIEM products because the combination of the SIEM that runs extremely efficiently on top of observability data that runs very efficiently on top of Flex Log, but also saves an immense amount of time by getting 90% of the issues out of the way with automated investigation that's extremely attractive to customers. All right. And I think with that, we're going to close the call. So -- before we go, I just want to give one quick shout out to the team because I know, as I said earlier, we have quite a lot going on in Q4, whether it's on the planning side, on the product building side or on the sales side, where I said we have a really, really exciting pipeline. And so we have a lot to do. I want to thank the team for the hard work there. I also -- I'm looking forward to meeting a lot of our existing and new customers at AWS re:Invent in a few weeks, and I'll see you all there. Thank you all.

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