

**Datadog, Inc. NasdaqGS:DDOG**  
**FQ2 2025 Earnings Call Transcripts**  
**Thursday, August 7, 2025 12:00 PM GMT**

# Table of Contents

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

# Call Participants

## EXECUTIVES

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

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

**Yuka Broderick**  
*Investor Relations*

## ANALYSTS

**Andrew Lodovico DeGasperi**  
*BNP Paribas Exane, Research Division*

**Brent John Thill**  
*Jefferies LLC, Research Division*

**Jacob Roberge**  
*William Blair & Company L.L.C.,  
Research Division*

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

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

**Mark Ronald Murphy**  
*JPMorgan Chase & Co, Research  
Division*

**Matthew Vincent Martino**  
*Goldman Sachs Group, Inc., Research  
Division*

**Michael Joseph Cikos**  
*Needham & Company, LLC, Research  
Division*

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

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

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

# Presentation

## Operator

Good day, and thank you for standing by. Welcome to the Q2 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 speaker today, Yuka Broderick, SVP of Investor Relations. Please go ahead.

## **Yuka Broderick** *Investor Relations*

Thank you, Didi. Good morning, and thank you for joining us to review Datadog's Second 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 third 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 March 31, 2025. Additional information will be made available through our upcoming Form 10-Q for the fiscal quarter ended June 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](https://investors.datadoghq.com).

With that, I'd like to turn the call over to Olivier.

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

Thanks, Yuka, and thank you all for joining us this morning to go through our results for Q2. Let me begin with this quarter's business drivers. Overall, we saw trends for usage growth from existing customers in Q2 that were higher than our expectations.

We experienced strong growth in our AI native cohort. The number of AI native customers are growing meaningfully with us as they see rapid usage growth with their products. Meanwhile, we saw consistent and steady usage growth in the rest of the business. We continue to see the overall demand environment as solid with an ongoing healthy pace of cloud migration and digital transformation, and 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 Q2 financial performance and key metrics, revenue was \$827 million, an increase of 28% year-over-year and above the high end of our guidance range. We ended Q2 with about 31,400 customers, up from about 28,700 a year ago. This includes about 150 new customers from our Eppo and MetaPlane acquisitions.

We ended Q2 with about 3,850 customers with an ARR of \$100,000 or more, up from about 3,390 a year ago, and these customers generated about 89% of our ARR, and we generated free cash flow of \$165 million with a free cash flow margin of 20%.

Turning to platform adoption. Our platform strategy continues to resonate in the market. At the end of Q2, 83% of customers were using two or more products, the same as last year, 52% of customers were using four or more products, up from 49% a year ago, 29% of our customers were using six or more products, up from 25% a year ago and 14% of our customers were using eight or more products, up from 11% a year ago.

So our customers continue to adopt more products, including our security offerings. As a reminder, our security customers can identify, manage vulnerabilities, with Code Security, Cloud Security, and Sensitive Data Scanner, and they can detect and protect from attacks with app and API protection, Workload Protection and Cloud SIEM.

We are pleased that our security suite of products now generates over \$100 million in ARR and is growing mid-40s percent year-over-year. While we are pleased to achieve this milestone, we're still just getting started in selling customer products in this area with new innovations such as our Bits AI Security Analyst.

Moving on to R&D. We held our DASH user conference in June, where we announced over 125 exciting new products and features for our users. So let's go through some of the announcements. First, we launched fully autonomous AI agents, including Bits AI SRE Agent to investigate alerts and coordinate incident response, Bits AI Dev Agent, an AI-powered coding assistant to proactively fix production issues and Bits AI Security Analyst to triage Datadog Cloud SIEM signals.

To further accelerate our users' incident response, we announced AI Voice Agent for Incident Response, so users can quickly get up to speed and start taking action on their phones. We also announced Handoff Notifications that make it easy to jump straight into the relevant context and quickly communicate with our responders and Status Pages to enable automatic updates for customers who undergo an incident.

Second, we delivered a series of products to help customers ship better software with confidence. With the Datadog Internal Developer Portal, developers can ship better and faster by gaining a real-time view into their software systems and APIs with the Software Catalog by provisioning infrastructure, scaffolding new services and managing code changes and deployments with Self-Service Actions and by following engineering and readiness standards with Scorecards.

We launched the Datadog MCP server to enable AI agents to access telemetry from Datadog and to act as a bridge between Datadog and MCP compatible AI agents like OpenAI Codex, Cursor and Claude Code by Anthropic. We work together with OpenAI to integrate our MCP server within the OpenAI Codex CLI, and the Datadog Cursor Extension now gives developers access to Datadog tools and observability data directly within the Cursor IDE.

Third, we are reimagining observability to meet our customers' increasingly complex needs. Our APM latency Investigator formulates and explores hypothesis in the background, helping teams to quickly isolate root causes and understand impact without combing through large amounts of data.

Proactive App Recommendations help users stay ahead of growing system complexity by analyzing APM data to detect issues and propose fixes before they become problems. We announced a Flex Frozen tier, so customers can keep logs in fully managed storage for up to 7 years and be able to search without data movement or rehydration.

Archived Search now enables teams to query archive logs directly in cloud storage like Amazon S3 buckets or in the Flex Frozen tier, and Datadog now supports advanced data analysis features within Notebooks.

Fourth, our security products cover new AI attack vectors across the application, model and data layers. At the AI data layer, Sensitive Data Scanner can now prevent the leakage of sensitive data and training data as well as LLM prompts and responses. At the model layer, we help secure against Supply Chain attacks in open source models and prevent Model Hijacking attacks. At the application layer, we help prevent Prompt Injection attacks and Data Poisoning in Runtime.

And finally, we showcased our new end-to-end AI and Data Observability capabilities. Engineers and machine learning teams can use GPU monitoring to gain visibility into GPU fleets across cloud, on-prem and GPU-as-a-service platforms such as CoreWeave and Lambda Labs.

With AI Agent Console, enterprises can monitor the behavior and interactions of any AI agent used by their teams. We now offer LLM Observability Experiments to help understand how changes to prompts, models or AI providers influence application outcomes.

We added a new Agentic flows visualization to LLM Observability to capture and understand the decision path of AI agent. And last but not least, and accelerated by our recent acquisitions of MetaPlane, Datadog now offers a complete approach to Data Observability across the entire data life cycle from iteration to transformation to downstream usage.

So we continue to relentlessly innovate to solve more problems for our customers. In doing so, we are being rightfully recognized by independent research, and we are pleased that for the fifth year in a row, Datadog has been named as a Leader in the 2025 Gartner Magic Quadrant for Observability platforms. We believe that this validates our approach to deliver a unified platform, which breaks down silos across teams.

Now let's move on to sales and marketing. We had a number of great new logo wins and customer expansions this quarter. So let's go through a few of those. First, we signed a 7-figure annualized expansion in a 3-year contract worth more than \$60 million with one of the world's largest banks. This company believes getting to the cloud is essential, so they can use AI on their extremely rich dataset to

improve how they manage risk and serve their customers. They are using Datadog as their strategic cloud observability platform, and they continue to migrate more applications to the cloud. This customer is expanding to 21 Datadog products with thousands of users who log into the Datadog platform every month.

Next, we signed a 7-figure expansion to an 8-figure annualized contract with a leading U.S. insurance company. Datadog is supporting this customer's efforts to consolidate observability tools and expand their cloud-based products. By adopting Datadog, they are experiencing fewer and less severe incidents with estimated savings of over \$9 million per year in incident response costs and improving more than 100,000 customer transactions that would otherwise be impacted every year. With this expansion, this customer will adopt 19 Datadog products and will consolidate a couple dozen tools across multiple business units.

Next, we signed a nearly 7-figure annualized expansion with a leading American media conglomerate. This customer has about 100 observability tools across more than 300 business units, and this tool fragmentation has resulted in inefficiencies, in extra costs and lost engineering time. They are expanding to 21 Datadog products, including all of our security products and replacing their paging solution with Datadog On-Call and Incident Management.

Next, we landed a 7-figure annualized deal with a leading Brazilian e-commerce company. This customer's previous observability vendor was unable to support them as they moved to newer software platforms and modern cloud infrastructure. By replacing this tool with Datadog, the company was able to gain full visibility into its cloud tech SaaS and saw significant improvements in application stability and incident resolution times. This customer will start with 7 Datadog products, including Flex Logs.

Next, we landed a 7-figure annualized deal with the delivery app of a major American retailer. This customer found our RUM and error tracking products to be immediately valuable, finding an issue on the first day of their Datadog trial that they hadn't identified after months of searching with their old tool. By adopting Datadog with 7 products to start, this customer will consolidate half a dozen tools while meeting their PCI compliance requirements.

Finally, we welcome back a leading U.S. mortgage company in a nearly 7-figure annualized deal. This customer had moved to using a dozen open source disconnected tools, which led to fragmented visibility, alert fatigue and poor customer experience. In returning to Datadog, they plan to adopt 6 products, including replacing their paging system with Datadog On-Call. And that's it for another productive quarter from our go-to-market teams who are now very hard at work on a busy Q3.

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. As we think about AI, we are incredibly excited about our opportunities. First, AI is a tailwind for Datadog as increased cloud consumption drives more usage of our platform. Today, we see this primarily in our AI native group of customers who are monitoring their cloud-native applications with us. There are hundreds of customers in this group. They include more than a dozen that are spending over \$1 million a year with us and more than 80 who are spending more than \$100,000, and they include 8 of the top 10 leading AI companies.

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 we continue to see rising customer interest for next-gen AI observability and analysis. Today, over 4,500 customers use one or more Datadog AI integrations.

Second, next-gen AI introduces new complexity and new observability challenges. Our AI observability products help our customers gain visibility and deploy with confidence across their entire AI stack, including GPU Monitoring, LLM Observability, AI Agent Observability and Data Observability, and we will, of course, keep innovating as the AI landscape develops further.

Third, we are incorporating AI into the Datadog platform to deliver more value to our customers. As I discussed earlier, we launched Bits AI SRE Agent, Dev Agent and Security Agent. We are seeing very good results with those with more improvements and new capabilities to come.

Finally, as a SaaS platform focused on our customers' critical workflows, we have a large volume of rich clean and detailed data, which allows us to conduct groundbreaking research. A great example of that is our TOTO foundational model for time series forecasting, which shows state-of-the-art performance on all benchmarks, even going well beyond specialized observability use cases, and you should expect to see more from us on that front in the future as well as taking novel research approaches and models straight into our products to improve customer outcomes.

So we are extremely excited about our progress so far against what we expect to be a generational growth opportunity. In other words, we're just getting started. And with that, I will turn it over to our CFO. David?

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

Thanks, Olivier. Q2 revenue was \$827 million, up 28% year-over-year and up 9% quarter-over-quarter. Now to dive into some of the drivers of this Q2 revenue growth. First, overall, we saw trends for usage growth from existing customers in Q2 that were higher than our expectations. This included strong growth in our AI native cohort as well as usage growth from the rest of the business that was consistent with recent quarters amidst a healthy and steady cloud migration environment.

We saw a continued rise in contribution from AI native customers in the quarter who represented about 11% of Q2 revenues, up from 8% of revenues in the last quarter and about 4% of revenues in the year ago quarter. The AI native customers contributed about 10 points of year-over-year revenue growth in Q2 versus about 6 points last quarter and about 2 points in the year ago quarter.

Now as previously discussed, we do see revenue concentration in this cohort in recent quarters. But if we look at our revenue without the largest customer in the AI native cohort, our year-over-year revenue growth in Q2 was stable relative to Q1.

We remain mindful that we may see volatility in our revenue growth on the backdrop of long-term volume growth from this cohort as customers renew with us on different terms and as they may choose to optimize cloud and observability usage over time.

As you heard from Oli, we continue to believe that adoption of AI will benefit Datadog in the long term, and we believe that the growth of this AI native customer group is an indication of the opportunity to come as AI is adopted more broadly and customers outside the AI native group begin to operate AI workloads in production.

Now regarding usage growth by customer segment. In Q2, our year-over-year usage growth was fairly similar across segments relative to previous quarters as SMB and mid-market usage growth improved in Q2, while enterprise customer usage growth remained roughly stable. Note that we are excluding the AI native cohort for the purposes of this commentary, and as a reminder, we define enterprise as customers with 5,000 or more employees, mid-market as customers with 1,000 to 5,000 employees and SMB as customers with less than 1,000 employees.

Regarding our retention metrics, our 12-month trailing net retention percentage was about 120%, higher than the high 110s last quarter, and our trailing 12-month gross revenue retention percentage remains in the mid-to-high 90s.

Now moving on to our financial results. First, billings were \$852 million, up 28% year-over-year, and remaining performance obligations, or RPO, was \$2.43 billion, up 35% year-over-year. Our current RPO growth was in the low 30s year-over-year, and our RPO duration was up slightly year-over-year.

As previously mentioned, we continue to believe that revenue is a better indication of our business trends than billings and RPO as those can fluctuate relative to revenue based on the timing of invoicing and the duration of customer contracts.

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 \$669 million for a gross margin of 80.9%. This compares to a gross margin of 80.3% last quarter and 82.1% in the year ago quarter.

As we've discussed in the last call, we saw an increasing impact of our engineers' cost savings efforts throughout this quarter as they delivered on cloud efficiency projects. And we are continuing our focus on cloud efficiency and believe that we have further opportunity for gross margin improvement in the second half of the year.

Our Q2 OpEx grew 36% year-over-year, up from 29% last quarter. As we've communicated over the past year, we plan to grow our investments to pursue our long-term growth opportunities, and this OpEx growth is an indication of our execution on our hiring plans. Q2 operating income was \$164 million for a 20% operating margin compared to 22% last quarter and 24% in the year ago quarter. Within that, as we've noted, we held our DASH user conference in June. And as expected, the event cost \$13 million. We also experienced a rising impact from the weaker dollar and absorbed \$6 million of negative FX impact during Q2. Excluding those expenses, operating income would have been 22% in Q2 or 200 basis points higher.

Now turning to the balance sheet and cash flow statements. We ended the quarter with \$3.9 billion in cash, cash equivalents and marketable securities, and our cash flow from operations was \$200 million in the quarter. After taking into consideration capital expenditures and capitalized software, free cash flow was \$165 million for a free cash flow margin of 20%. Now for our outlook for the third quarter and the remainder of fiscal 2025.

First, our guidance philosophy overall remains unchanged. As a reminder, we base our guidance on recent trends observed and apply conservatism on these growth trends. For the third quarter, we expect revenues to be in the range of \$847 million to \$851 million, which represents a 23% year-over-year growth.

Non-GAAP operating income is expected to be in the range of \$176 million to \$180 million, which implies an operating margin of 21%, and non-GAAP net income per share is expected to be \$0.44 to \$0.46 per share based on approximately 364 million weighted average diluted shares outstanding.

For fiscal 2025, we expect revenue to be in the range of \$3.312 billion to \$3.322 billion, which represents a 23% to 24% year-over-year growth. Non-GAAP operating income is expected to be in the range of \$684 million to \$694 million, which implies an operating margin of 21%, and non-GAAP net income per share is expected to be in the range of \$1.80 to \$1.83 per share based on approximately \$364 million average diluted shares.

Some additional notes on our guidance. We expect net interest and other income for fiscal 2025 to be approximately \$150 million, and due to the impact of the recent federal tax legislation, we now expect cash taxes for 2025 to be about \$10 to 20 million. 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% to 5% of revenues in fiscal year 2025.

To summarize, we are pleased with our execution in Q2, including the many products and features we launched at DASH. We are well positioned to help our existing and prospective customers with their cloud migration and digital transformation journeys, including their adoption of AI. I want to thank all Datadogs worldwide for their efforts. And with that, we'll open the call for questions. Operator, let's begin our Q&A.

# Question and Answer

## Operator

[Operator Instructions]

And our first question comes from Raimo Lenschow of Barclays.

### **Raimo Lenschow**

*Barclays Bank PLC, Research Division*

Perfect. Two quick questions from me. Olivier, like you talked about the AI contribution and slowly broadening out. How should we think about it in terms of when this goes much broader into inference, et cetera? So does that mean everyone like Barclays, JPMorgan, et cetera, they all kind of need to do more around observability because they're going to do more inference, et cetera.

So in a way, like OpenAI, et cetera, is just setting the scene for future? And what do you think about the market opportunity there? And then, David, in the second half of last year, you hired a lot of extra sales guys. Can you talk a little bit about that ramp and where they are in their productivity curve?

### **Olivier Pomel**

*Co-Founder, CEO & Director*

Yes. On the AI opportunity, so there's really multiple layers to it. The first layer is largely what we see today, which is, companies that are running their inference stack and the application around it, in cloud environments. So that's the case of the model makers or if you think of the companies that are doing coding agents, things like that. That is what we see today, and it looks a lot like normal compute. So you have normal machine CPUs, some GPUs, quite a few other components, databases, web servers, things like that. So that's the bulk of what we see today. And there's going to be more of it as the AI applications come into production. There are more specialized inference workloads and even training workloads in some situations that rely on instrumenting GPUs. And for that, we have a new product out there that does GPU Monitoring that we announced at DASH. But all that I would call the infrastructure layer of AI.

Then on top of that, there's new problems in terms of understanding what the applications themselves are doing and the applications are largely nondeterministic anymore. They either are run by a model that is nondeterministic by nature or they run in code that was not as carefully written as it used to be.

It's not completely written by humans, just largely written by AI agents, and as a result, you also need to spend a lot more time understanding how that code is working and that largely happens in production. So that's a brand-new area of observability, which is how do you deal with applications that have not been completely defined in development and that have to be evaluated in production.

And what we think is the whole market is going there, not just the AI natives, the AI natives are definitely doing that today, both applications are running on models and code that has been largely written by agents, but the rest of the market is going there, and the best proof point you see of that is the very, very broad adoption today, both of the API-gated AI models and of the coding agents, which you see in every single large enterprise today.

### **David M. Obstler**

*Chief Financial Officer*

Yes. And as to sales capacity, we have been successful in increasing both our number of salespeople and our ramp sales capacity. We started that, as you said, in the last part of 2024, and we are seeing evidence of that through our new logo production and our pipeline. We need to, as we talked about previously, go through the ramping of that, but in looking at the size and productivity and performance, we see some good signs that, that core capacity is becoming productive.

## Operator

Our next question comes from Sanjit Singh of Morgan Stanley.

### **Sanjit Kumar Singh**

*Morgan Stanley, Research Division*

Congrats on the really stellar results this quarter. David, when I look at the guide, I mean, this is probably one of the more impressive guides coming out of Q2 that I've seen in a couple of years. If I square that against the commentary that you guys made on the AI native cohort that, look, there could be volatility from this cohort.

When I try to put those two together, the guidance is really strong, and so when I think about that potential risk, is it fair to assume that it's not something that you're seeing right now and may come to play later on down the road because the guidance seems really strong. It doesn't seem to -- at least on the face, doesn't seem to anticipate that much volatility from the AI native cohort.

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

Yes. I think we gave metrics indicating that based on what we saw in the quarter and we're seeing now that the AI cohort continues to grow quite rapidly, and we're winning a good market share in that, and so how we incorporate that into the guidance is, as we discussed previously, we know that there might be volatility in usage or in -- as we negotiate contracts in unit rates, and so therefore, we adopt conservative assumptions as to that performance in the remainder of the year. It's not something, as you can tell from the growth metrics that we see yet in our results, but as we learned in the previous cycle with cloud natives, there can be volatility, and we want to make sure we incorporate that in our guidance.

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

Perfect. And then, Olivier, with the new security disclosures, congrats on crossing the \$100 million threshold. Is there any sort of change in the buying behavior? There's been consolidation in the industry. You guys have been advancing your portfolio quite significantly. You guys have fully autonomous security agents. What's your prospect for this pool of the business, this part of the business to drive growth for the balance of the year and going into 2026?

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

Yes. So we have a very good product set, and we mentioned we have a few different products in there. There are a couple of those products that are really, I would say, reaching an inflection point in terms of what they're doing on the customer side.

When I think of where we're successful today in security, we're very successful at getting broad adoption, like a large number of customers, a few customers that are spending \$1 million plus on security with us. So we're good with the -- we're very happy with the proof points we have there.

What we haven't done very well yet is getting standardized adoption wall-to-wall in large enterprises, and that's the next focus for us on the security side, and some of that is product work, but a lot of it is a few customizations to go-to-market there, so we get better at selling enterprise-wide security top down, which is not something we have done a lot of in the past. So that's sort of where we are as a product.

So happy with where we are. A lot of groundwork has been done on the product side, but there's quite a bit more work to be done and a ton more opportunity in front of us. So we are -- that's why we're focusing on it.

**Operator**

And our next question comes from Kash Rangan of Goldman Sachs.

**Matthew Vincent Martino**  
*Goldman Sachs Group, Inc., Research Division*

This is Matt Martino on for Kash Rangan. David, you called out enterprise consumption volatility last quarter. It sounds like that may have been consistent this time around while SMB continues to improve. So could you perhaps characterize any discernible trends between these two customer demographics? What went right relative to your expectations heading into 2Q and really, how that informs your second half guide?

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

Yes. I think broadly, we're calling out that the usage trends across the segments were roughly consistent with the previous quarters. We said we did see some more concentrated - this is not a comment about AI, this is a comment about enterprise - Enterprise took down less

consumption relative to a spike, but we saw that stabilize, and we've seen small, but gradual improvement of the SMB as a result of their usage of our products.

**Operator**

And our next question comes from Mark Murphy of JPMorgan.

**Mark Ronald Murphy**

*JPMorgan Chase & Co, Research Division*

Congrats. So Olivier, I actually wanted to ask you about Toto and BOOM, those announcements. It looks like you're bringing very serious AI research to a space where it is applicable and opening it up very broadly, the size of the dataset is vast. I'm curious what type of response do you expect to see here? And just help us understand maybe how that can sustain growth in future years? And then I have a quick follow-up for David.

**Olivier Pomel**

*Co-Founder, CEO & Director*

Look, we think there's so much opportunity in automation with autonomous AI agents. Like we really broke it out in three different categories so far. One is the SRE and responding to alerts and investigating alerts and remediating those issues.

Second one is coding, fixing issues that we find in the code that happened in production and verifying these fixes ourselves. And the last one is security, investigating security signals on our own so that customers don't have to do that themselves.

There's so much I missed and that can happen there. A lot of it is going to depend on great research, which is why we built a research team and which is why we developed and released an open weight research model already. Of course, the next step after releasing these research models is to incorporate them into the product. So that's also one of the things we're working on right now, but there's just so much opportunity in front of us there that we're -- at this point, we're happy we got a great start.

We got fantastic results in our first release, research output is really like a state-of-the-art model that beats every single other model in a category that has seen quite a bit of action over the years. Time series forecasting has very wide applicability in a lot of different domains. So I think it shows that we can perform at the highest level there, and I think it's a great sign of things to come in terms of AI automation and AI agents.

**Mark Ronald Murphy**

*JPMorgan Chase & Co, Research Division*

Okay. And then, David, we keep pointing out that Datadog is one of the only software companies that's investing seriously in headcount growth and it feels like that is paying top line dividends pretty tremendously today. We noticed the R&D spending is up noticeably in Q2.

Just wondering what are the mechanics that are driving that on the R&D line? And then the flip side is what's allowing you to guide operating income so much higher in Q3 than you had guided that for Q2?

**David M. Obstler**

*Chief Financial Officer*

Yes. In R&D, as we talked about, we had an aggressive investment plan, and we've been able to execute, and I think our recruitment -- credit to our recruitment team, we've been able to get people in the door, the right people earlier in the year.

There are some things within that around FX that weigh a little bit on it because, as you know, we do have a significant R&D center in Paris, but I think the overall trend is the execution and recruiting.

We talked about some of the factors in Q2 that caused the operating expenses to increase at a rate of 36%, and some of those are things like the timing of DASH. We talked about \$13 million, the FX, and I think that we have good line of sight on the drivers in R&D, both in terms of -- as we talked about and some of the operating expenses are -- have some seasonality in it.

**Olivier Pomel**

*Co-Founder, CEO & Director*

The one thing I would add, which is that we also are spending more on AI training and inference in R&D if you compare to past years in R&D and the output of that is things such as Toto or the next versions of it that we're training right now and experiments

we're running to train agents, simulations to train agents and things like that. You shouldn't expect the overall picture of our R&D investment to change in the future, although I think we expect the same envelope to be what we use moving forward.

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

Yes. I'll add that and really call out to our R&D team and our FinOps that we said last quarter that we were going to focus on how we use cloud. That applies to the gross margin. And as you know, we dog-food. We use a lot of our applications internally, and we were quite successful in Q2. That run rate, we expect to continue forward in optimizing our cloud usage, which will have an effect on the margins and the OpEx growth rates as we proceed through the year.

**Operator**

And our next question comes from Koji Ikeda of Bank of America..

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

We all see that the second quarter was really, really strong. Guidance for 2025 looks really, really great. And so I wanted to ask you about contract visibility. How are you feeling about contract visibility, specifically with your large AI native customers? I have to imagine you're very close to these customers and having lots of conversations with them. And so I know there is some concern about there. And David, you mentioned potential volatility. So I really want to ask about how you're feeling about contract visibility.

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

I mean, look, we can't really speak about any specific customers. As a reminder, any individual customer can do whatever they want. They are the heroes of their own stories, and we can't really speak for them. I would say we have strong product engagement from our top customers in general. We are working on making it -- making Datadog the very best platform for every company at any scale, including scales that have never been seen before in companies with high growth, and I would say that's about it.

When you look at the way we forecast the business, remember that we have overall an extremely high retention product. For most customers, it's not rational to do it themselves, build their own solutions. We have many customers who did churn to build themselves, who come back afterwards, and we named one on the call today. So we feel confident about the way we forecast the business and the mid-to-long term there.

Of course, as we renegotiate with customers, as they increase volume, et cetera, et cetera, what typically happens is we see short-term drops and long-term growth in the revenue associated with them, and that's the way we've always operated.

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

And I did have a follow-up on security, and so it sounds -- I mean, great to hear about the milestones, \$100 million, growing 40%, and so thinking about the product set, how are you thinking about expanding the capabilities from here? Are you focused on more organic, inorganic? And maybe an update to your M&A philosophy. I mean, I guess the question here is, are you willing to go much bigger to supplement your security strategy?

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

Look, we're looking at a number of different things in security; there's a lot of companies out there. There's a lot of product areas we cover already and a lot of more product areas we can cover. It's also a space where you need to cover a lot of the, how do we call them, boring must-have table stakes features on one end, but also there's quite a bit of investment in the future with the way the whole field is being disrupted with AI. So there's quite a bit of work to be done there. You should expect us to do more M&A around that as we do in the rest of the business as there is a lot of assets out there, and there's a lot of opportunities to grow.

**Operator**

And our next question comes from Karl Keirstead of UBS.

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

Okay. Great. Maybe I'll direct this to David and link the AI native exposure to margins. So David, now that the AI natives are 11% of Datadog's revenue mix, I think it's fair to ask whether the revenues from that cohort are coming at similar margins as the rest of the business? Or do you think that this could be even short term, a modest source of margin pressure?

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

Yes. I would say like we talked about last quarter, this isn't about the AI and margins, the AI cohort versus non-AI cohorts. We price based on volume and on term. So to the extent you would have an AI customer who's doing much the same things as our other customers in the use of the product, has similar volumes and similar terms to the non-AI, it would be similar margins.

To the extent that we have a larger customer in there, given our price grids, that customer would get a better discount. That's the way we've always priced. So it really is related to customer size rather than AI native or non-AI native.

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

And with a bit of it in commercial, so we did see, as we mentioned last quarter, we were seeing gross margins going down a little bit further than we would like them to. So what happened is we tasked our engineering teams with optimizing the cloud usage, which goes across all of our customer base.

What we did is, we turned to our own product, we turned to our Cloud Cost Management product and our Profiling product largely, and then we, in a matter of months, will really turn up substantial improvements, savings on our bills and improvements in performance and efficiency of our systems, while we're still shipping new features, and that's something that we're working right now to bring to all of our customers so they can get the same effect and they can see their margins go up as well.

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

Got it. And maybe the natural follow-up there is, David, you mentioned that you're optimistic about gross margins in the second half. Is that because of what Olivier just mentioned? Or are there some other drivers you have in mind?

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

No, it's because of what Olivier mentioned. So we said we were engaging in these efforts. And as we were more successful in the quarter, we will be carrying that run rate forward, which wasn't fully in Q2 as well as using what Olivier mentioned, using Cloud Cost Management and our projects to have further opportunities going forward. So it's really about our progress and pace, which has been successful in our cloud efficiency going forward.

**Operator**

And our next question comes from Mike Cikos from Needham.

**Michael Joseph Cikos**  
*Needham & Company, LLC, Research Division*

I just wanted to double back on the enterprise segment, and just -- this is for Oli. But if I'm thinking about it, I know that we have the enterprise demonstrating the stable growth. Is it fair to assume -- like is the analogy for enterprises who are more traditionally using CPU versus the AI native companies who are growing investment in GPUs? Is it analogous to like 15 years ago where we saw, hey, on-prem continues to see investment, but maybe more dollars are going towards cloud. Is that like a fair analogy when we think about what sort of behavior is exhibited by these different customers and where Datadog is headed?

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

I don't know if you can say it exactly this way because at the time, the on-prem versus cloud has tended to be the same customers, whereas today, like the AI natives and the enterprise are different companies altogether.

I think the main difference is the AI natives have businesses that are growing very, very fast and infrastructure that are growing very, very fast themselves, whereas the enterprises are still going through a controlled migration from on-prem into the cloud, and the

rate there is more limited by their bandwidth to undergo that migration as opposed to being driven by an explosion of traffic on the demand side for them.

If I look at our enterprise segment in general, we see great trends in terms of the bookings, in terms of new products attached, new customers, things that these customers are buying from us that are net new, but we see that the usage growth is a bit more moderate than that at this point, and I think that speaks to the bandwidth on their end just to move the workload and to go fast there.

And that relates in part to the fact that a lot of the attention is spent on figuring out what AI technologies they're going to adopt and how they're going to ship these AI applications into production.

Overall, we see that rate as stable. So we think this is healthy, but we think this is -- we think we will see more growth from these enterprise customers as they actually get into production with the AI applications in the future.

**Michael Joseph Cikos**  
*Needham & Company, LLC, Research Division*

Understood, and congrats on the security. I didn't want to leave hanging. I don't know if we got commentary on it, but could we please get an update on Flex Logs? I know it was a shining star if I go back a quarter ago, but just wanted to see how progress is tracking on the Flex Log side of the house.

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

Yes. All of the big deals with enterprise customers now involve Flex Logs in some form, and that's a story that resonates very well when we -- especially when we have customers that want to migrate from legacy solutions from logs. So there's a number of things that we're working on with them, in particular, making sure the migration is painless for them, and there's a number of things that we are investing in on that side.

But Flex Logs is a big draw for them as it really changes the picture economically and the predictability of the observability cost for them, which is a major concern for data-intensive observability such as logs.

**Operator**

And our next question comes from Jake Roberge of William Blair.

**Jacob Roberge**  
*William Blair & Company L.L.C., Research Division*

There's obviously been a lot of talk about AI natives around the business. I know you've talked about the potential for optimization for several quarters, but we continue to see really strong growth in that segment. So if you were to see optimization, when would you expect that to happen? And as you get a wider swath of customers in that AI native cohort, do you think you're at the place where you could actually digest an optimization by one or two of those customers?

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

Well, I mean, look, if I knew when it was going to happen, I would tell you. The nature of our customers is they grow, they have their own businesses to run. They have their own constraints. We're here to help them deliver their services, and that's what we work on every single day. Now every now and then, there's a renegotiation, a renewal on occasions for customers to figure out what they need to optimize and what they need to do for the future. But we never know whether it's going to happen this quarter, next quarter, in three quarters next year, never. That's really hard to tell.

**Jacob Roberge**  
*William Blair & Company L.L.C., Research Division*

Okay. That's helpful, and then could you also talk about the uptake and feedback that you're getting for your own AI solutions like Bits AI, the new observability agents? And when do you think those could really start layering into the model?

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

Yes. So I mean the initial response to the AI agents is really pretty positive. So the AI's actually works surprisingly well. I mean if you think of how far the technology has grown in a number of a couple of years, and so right now, we're busy basically shipping it to as many customers as we can and enabling the customers with it, and that's a big area of focus in the business as well.

I think it was developed by a fairly small team, the actual product that we ship, and now we're busy scaling that up as fast as we can so we can serve all those customers. That's the core focus of the business today. So the initial response is very positive. We've had customers purchase it pretty quickly in their trials, and so we feel very good about it.

**Operator**

And our next question comes from Brent Thill of Jefferies.

**Brent John Thill**

*Jefferies LLC, Research Division*

David, just on the quota-carrying rep capacity, I know you've been investing aggressively ahead of the curve. But when you think about 2025, are you accelerating that count based on the great results you've seen? Are you digesting that count given those reps are on board? Just give us a sense and flavor of what that quota rep count looks like through the rest of the year, and if you can shape the year how that looks versus '24.

**David M. Obstler**

*Chief Financial Officer*

Yes. What we're doing is we're executing the plan we entered the year with. We knew -- I think we said that we had underinvested in go-to-market and looked at that with the white space, et cetera, and I would say we're successfully executing that. The plan was a little more front weighted given our appetite for taking advantage of that opportunity, but we're executing that, and we will look at the -- towards the end of the year as we plan for next year on the metrics around that and try to calibrate how we look at that growth next year.

**Brent John Thill**

*Jefferies LLC, Research Division*

Okay. And Olivier, I'm just curious, many CEOs are either holding headcount flat or down. We've seen Meta headcount down from 2 years ago, Microsoft headcount flat, others -- Palantir saying they're going to shrink headcount and 10x revenue. Do you believe you can become more efficient with fewer? Or do you think that, that model doesn't apply that you're seeing with other software companies?

**Olivier Pomel**

*Co-Founder, CEO & Director*

I mean, look, there's definitely -- the spend is shifting a little bit on the engineering side. As I said, we compute -- we consume more AI training inference, and so that's definitely changing a bit of the balance between what you have humans do and what you offload to GPUs.

That being said, we're still completely constrained by the amount of product we can put out there. There's a ton of opportunity in every single direction we look, whether that's on the AI automation, whether it's on the security side, whether that's in the new areas, just better observability or experimentation that we're going after, and so for us, there's very strong ROI in the adds that we're making at the moment.

**Operator**

And our next question comes from Andrew DeGasperi of BNP Paribas.

**Andrew Lodovico DeGasperi**

*BNP Paribas Exane, Research Division*

First, on the ramp-up in terms of sales capacity, would you say that's been broad-based in terms of the productivity across both international and domestic?

**David M. Obstler**

*Chief Financial Officer*

As we talked about previously, we have a less developed international footprint, and so our growth rate internationally is running higher. We have markets we've talked about before like Brazil and India and parts of APJ and Middle East that we have opportunities to grow our footprint. So we are executing in that way.

We're doing a bottoms-up as always. We're looking at the accounts. We're looking at the TAM, and we're looking at how much we're covering it. So that produces a result of a little more investment intensity internationally versus in North America, but there are lots of opportunities in North America as well.

**Andrew Lodovico DeGasperi**  
*BNP Paribas Exane, Research Division*

That's helpful, and then on the enterprise side, I mean, given some of these reps are obviously on the ground, should we expect the number of the attach rates in terms of the 3 or 4 more products per customer sort of accelerate at this level? I know they've been ticking up about 1 point every quarter. Just wondering if that's something we should be seeing.

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

Well, I think broadly, we expect the trends that we've seen of landing with some of the core products in the pillars and then expanding to continue. As the platform has expanded, we've tended to land with more products, but those trends that we evidenced in the script, we expect to continue in the geographies.

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

And keep in mind, a lot of the -- so when you're in the field, it's always easy to upsell a customer than to land a new customer, and a lot of the work we're doing in territory management and in co-planning for the sales team is really to make sure that there's enough of an incentive to go and look for new customers. So we keep driving the number of new customers up as well. So there's this balance always between whether you direct the sales force at upselling existing customers or landing new customers.

## **Operator**

And our next question comes from Patrick Colville of Scotiabank.

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

And I guess I just wanted to say before I ask my question, congrats on the S&P 500 Index inclusion. I mean that's a really nice milestone for you guys. Look, the question we get consistently from investors is on competition. I mean you referred to your views on competition kind of tangentially in other kind of answers, but maybe more specifically, I mean, what are you seeing competitively in observability? And the one we get asked about a lot is versus Grafana and Chronosphere.

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

Yes. I mean, look, there's always been competition in the field. As I like to say, when I first fundraised for Datadog, the world that was coming back to me every single time with every single node I was getting from our VCs was crowded space, and so throughout the life of the company, there's been not only incumbents that we mostly have been in the market now, but also a steady stream of new entrants that we also have year after year, have been in the market.

There's always new companies, always new folks that are building new things in observability. I think it's very attractive for engineers to build that. I would know something about it. Generally speaking, the competitive landscape hasn't changed much in the past 10 to 15 years, about the same.

The way we win and we will keep winning is by offering an integrated platform that solves as many problems as possible for our customers end-to-end. So we solve -- we don't just focus on one. We don't just focus on one data store, one specific brick that our customers might want to use. We solve the whole problem for them end-to-end.

And then in the long run, we win by being more innovative, by having an economic model that lets us invest more in R&D, develop more products, build the existing products into the future faster than anybody else can do and cover more adjacencies faster than

anybody else can do so we can have the broadest platform. So that's the reason we win. And if you look at all of the companies you mentioned, none of them are in a position to do the same. And so that's where we're going to end up winning in the end.

And I think that's at the end of the call. So that would be the last question. And just to close out, I want to thank our customers for working with us to bring all of those great new products to market. So we had a lot on our plate this year. You've seen that at DASH.

It was amazing, by the way, to see all these customers and meet with them at DASH and see the reception we get for all these new products. And so I want to thank them. And I know we're working with many of them on how these products are going to be adopted and what's going to happen in Q3 and Q4. So again, thank you, and I will see you next quarter.

**Operator**

This concludes today's conference call. Thank you for participating, and you may now disconnect.

Copyright © 2025 by S&P Global Market Intelligence, a division of S&P Global Inc. All rights reserved.

These materials have been prepared solely for information purposes based upon information generally available to the public and from sources believed to be reliable. No content (including index data, ratings, credit-related analyses and data, research, model, software or other application or output therefrom) or any part thereof (Content) may be modified, reverse engineered, reproduced or distributed in any form by any means, or stored in a database or retrieval system, without the prior written permission of S&P Global Market Intelligence or its affiliates (collectively, S&P Global). The Content shall not be used for any unlawful or unauthorized purposes. S&P Global and any third-party providers, (collectively S&P Global Parties) do not guarantee the accuracy, completeness, timeliness or availability of the Content. S&P Global Parties are not responsible for any errors or omissions, regardless of the cause, for the results obtained from the use of the Content. THE CONTENT IS PROVIDED ON "AS IS" BASIS. S&P GLOBAL PARTIES DISCLAIM ANY AND ALL EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE OR USE, FREEDOM FROM BUGS, SOFTWARE ERRORS OR DEFECTS, THAT THE CONTENT'S FUNCTIONING WILL BE UNINTERRUPTED OR THAT THE CONTENT WILL OPERATE WITH ANY SOFTWARE OR HARDWARE CONFIGURATION. In no event shall S&P Global Parties be liable to any party for any direct, indirect, incidental, exemplary, compensatory, punitive, special or consequential damages, costs, expenses, legal fees, or losses (including, without limitation, lost income or lost profits and opportunity costs or losses caused by negligence) in connection with any use of the Content even if advised of the possibility of such damages. S&P Global Market Intelligence's opinions, quotes and credit-related and other analyses are statements of opinion as of the date they are expressed and not statements of fact or recommendations to purchase, hold, or sell any securities or to make any investment decisions, and do not address the suitability of any security. S&P Global Market Intelligence may provide index data. Direct investment in an index is not possible. Exposure to an asset class represented by an index is available through investable instruments based on that index. S&P Global Market Intelligence assumes no obligation to update the Content following publication in any form or format. The Content should not be relied on and is not a substitute for the skill, judgment and experience of the user, its management, employees, advisors and/or clients when making investment and other business decisions. S&P Global Market Intelligence does not act as a fiduciary or an investment advisor except where registered as such. S&P Global keeps certain activities of its divisions separate from each other in order to preserve the independence and objectivity of their respective activities. As a result, certain divisions of S&P Global may have information that is not available to other S&P Global divisions. S&P Global has established policies and procedures to maintain the confidentiality of certain nonpublic information received in connection with each analytical process.

S&P Global may receive compensation for its ratings and certain analyses, normally from issuers or underwriters of securities or from obligors. S&P Global reserves the right to disseminate its opinions and analyses. S&P Global's public ratings and analyses are made available on its Web sites, [www.standardandpoors.com](http://www.standardandpoors.com) (free of charge), and [www.ratingsdirect.com](http://www.ratingsdirect.com) and [www.globalcreditportal.com](http://www.globalcreditportal.com) (subscription), and may be distributed through other means, including via S&P Global publications and third-party redistributors. Additional information about our ratings fees is available at [www.standardandpoors.com/usratingsfees](http://www.standardandpoors.com/usratingsfees).  
© 2025 S&P Global Market Intelligence.