Datadog, Inc. NasdaqGS:DDOG

FQ2 2026 Earnings Call Transcripts

Thursday, August 6, 2026 8:00AM ET

Table of Contents

Call Participants

EXECUTIVES

Olivier Pomel
Co-Founder, CEO & Director

David M. Obstler
Chief Financial Officer

Yuka Broderick
SVP of Strategic Finance & Investor Relations

ANALYSTS

Gabriela Borges
Analyst, Goldman Sachs & Co. LLC

Samik Chatterjee
Analyst, JPMorgan Securities LLC

Mike Cikos
Analyst, Needham & Co. LLC

Andrew DeGasperi
Analyst, BNP Paribas Securities Corp.

Koji Ikeda
Analyst, BofA Securities, Inc.

Ritu Kashyap
Analyst, KeyBanc Capital Markets, Inc.

Ittai Kidron
Analyst, Oppenheimer & Co., Inc.

Raimo Lenschow
Analyst, Barclays Capital, Inc.

Howard Ma
Analyst, Guggenheim Securities LLC

Brad Reback
Analyst, Stifel, Nicolaus & Co., Inc.

Andrew Sherman
Analyst, TD Securities (USA) LLC

Sanjit K. Singh
Analyst, Morgan Stanley & Co. LLC

Alex Zukin
Analyst, Wolfe Research LLC

Presentation

Operator
Good day and thank you for standing by. Welcome to the Q2 2026 Datadog Earnings Conference Call. At this time, all participants are in a listen-only mode. After the speakers' presentation, there will be a question-and-answer session. [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
SVP of Strategic Finance & Investor Relations

Thank you, Lauren. Good morning and thank you for joining us to review Datadog's second quarter 2026 financial results, which we announced in our press release issued this morning. Joining me on the call today are Olivier Pomel, Datadog's Co-Founder and CEO, and David Obstler, Datadog's CFO.

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 2026, and related notes and assumptions, our product capabilities, and our ability to capitalize on market opportunities.

The words anticipate, believe, continue, estimate, expect, intend, will, and similar expressions are intended to identify forward-looking statements or similar indications of future expectations. These statements reflect our views 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, 2026. Additional information will be made available in our upcoming Form 10-Q for the fiscal quarter ending June 30, 2026, and other filings with the SEC.

This information is also available on the Investor Relations section of our website, along with a replay of this call. We will discuss non-GAAP financial measures, which are reconciled to their most directly comparable GAAP financial measures in 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
Co-Founder, CEO & Director

Thanks, Yuka. And thank you all for joining us to go over our Q2 results. Let me begin with this quarter's business drivers. Our revenue growth in Q2 has accelerated across our customer base.

On one hand, our AI-native customer cohort continued to grow and diversify, both in the number of customers we serve and the scale of those customers. But on the other hand, and as a great illustration of the breadth of strength across our business, revenue growth for our non-AI customers also accelerated again this quarter to the high-20s% year-over-year, up from the mid-20s last quarter and 18% in the year-ago quarter.

Overall, we continue to see healthy trends in customer demand. Our broad base of customers, from the most nimble startups to the largest and most established enterprises, are all adopting AI. We think this is accelerating their usage of cloud and modern technologies, as well as their usage of the Datadog platform to observe, secure and act on their cloud and AI workloads.

Regarding our Q2 financial performance and key metrics, revenue was $1.12 billion, an increase of 36% year-over-year and above the high end of our guidance range. We ended Q2 with about 33,400 customers, up from about 31,400 a year ago. We also ended with about 4,720 customers with an ARR of $100,000 or more, up from about 3,850 a year ago. These customers generated about 91% of our ARR.

And we generated free cash flow of $279 million, with a free cash flow margin of 25%.

Turning to product adoption. Our platform strategy continues to resonate in the market. For example, 58% of our customers now use four or more products, up from 52% a year ago. 37% of our customers use six or more products, up from 29% a year ago. And 13% of our customers use ten or more products, up from 7% a year ago. So we're landing more customers and delivering value across more products.

And our products are broadly delivering strong growth in usage and ARR. As an example, RUM, or real user monitoring, now exceeds $200 million in ARR and accelerated as it scaled to over 50% growth year-over-year. Our customers are sending more user sessions and using RUM in conjunction with our newer Product Analytics to optimize their business outcomes.

Moving on to R&D, we held our DASH user conference in June, where we announced over 100 exciting new products and features for our users. So let's go through some of the announcements.

First, we expanded Bits AI to accelerate and automate the DevOps loop. This is the loop that goes from detection to investigation to remediation that engineers go through each time something breaks.

At DASH, we announced a lot of new Bits capabilities for the DevOps loop. Bits can now create and maintain monitors, identify root causes within minutes of a negative signal, recommend and implement fixes, follow guardrails to add safety and controls, continuously learn and improve from prior incidents, and detect symptomatic behaviors early to repair infrastructure issues before they escalate.

Second, we announced Bits AI products to address the development loop. This is the loop that goes from coding to delivery to evaluation that developers navigate to get code to production. For this loop, Bits Release now acts as an AI release validation agent, analyzing the impact of code changes, running end-to-end checks, and verifying production rollouts. Bits Code generates code fixes, grounding every fix in real production behavior, and Bits Testing also automates synthetic test generation and maintenance.

Third, we expanded Datadog for AI, our products that observe, secure, and optimize the AI stack from end-to-end. Data Observability enables companies to trust the data being used by AI with lineage, quality monitoring and jobs monitoring. Bits Data Analysis uses our rich data context to accurately answer business questions, and Agent Console provides visibility into AI agent usage, cost, and effectiveness.

In Agent Observability, our Patterns capability automatically clusters user interactions into behavior groups to identify quality or cost issues, and Bits Evals handles the repetitive parts of the agent development loop in order to improve the outcomes of agents.

Fourth, we are broadening our platform to ingest, correlate, analyze, and act on more data, whether on-prem or in the cloud. In Network Monitoring, we launched Network Path and Network Configuration Management to trace changes that cause complex network issues. Within Database Monitoring, Bits Database Optimizer now automatically simulates and evaluates the impact of AI-generated changes in order to optimize slow queries.

In Log Management, Federated Logs enable users to query external data stores, including Databricks and ClickHouse. And with Bring Your Own Cloud, or BYOC, customers can now use the full Datadog experience on logs that are kept within their own infrastructure. And we've also announced that we are bringing BYOC to Metrics and Traces as well.

In the Digital Experience space, Journey Monitoring automatically gives a single shared view for every critical user flow. And for custom metrics data, we introduced Infinite Cardinality Metrics, which allow our users to answer arbitrarily complex questions as they generate larger amounts of data with AI agents without incurring any extra costs.

Finally, we launched a number of innovations to secure the AI stack and defend against a new class of AI-powered attacks. AI Guard Agent Discovery finds and maps every known and unknown custom agent, so security teams can see what is protected and what is not. AI Guard for Custom Agents provides runtime protections to block attacks that can only be detected with real-time observability data. AI Guard for Coding Agents applies the same deep observability to block malicious skills and packages in code.

And we also announced Runtime Prioritization Engine to cut vulnerability noise by over 95%. And finally, we expanded Bits Security Analyst to run on non-Datadog SIEMs, so customers can benefit from the smarts and the learnings of our broad data set, regardless of which SIEM they deploy.

As we continue to innovate, we are being rightfully recognized by independent research. We are pleased to see that for the sixth year in a row, Datadog has been named a Leader in the 2026 Gartner Magic Quadrant for Observability Platforms.

Let's move on to Sales and Marketing and look at a few of the deals our GTM teams have closed in what has been a very strong quarter. First, we landed a six-figure annualized deal with a Fortune 10 company. This company is expanding its e-commerce business, and they plan to use Datadog Log Management alongside 10 other Datadog products to improve customer experience and business outcomes. This win validates our expanded go-to-market approach to focus on the world's largest companies and win opportunities in the most complex environments.

Next, we landed seven-figure annualized deals with two neolabs. These AI labs are rapidly scaling their AI model training workloads and preparing for major product launches. By deploying observability using Datadog, they gain visibility across their training infrastructure and GPU fleets and can iterate faster on their AI models. They are also using Bits AI to rapidly build monitors, dashboards, and alerts for deep observability context.

Next, we landed a seven-figure annualized deal with a South American bank. This bank's fragmented legacy monitoring stack and manual triaging caused significant application downtime that was often called in by customers. By consolidating into Datadog with 11 products, this customer enables visibility from their mainframe all the way to their microservices, and has already reduced mean-time-to-resolution on live production incidents. They are adopting Cloud SIEM and Data Security and evaluating other Datadog Security products to improve their security posture.

Next, we signed a seven-figure annualized expansion for an eight-figure annualized deal with a Fortune 100 health insurance company. This customer's biggest pain point is to deliver great experience to their members throughout their care, while protecting PII across dozens of business units. Datadog's HIPAA compliance and PII handling in RUM, Log Management, and Cloud SIEM allowed us to differentiate and win over competitive solutions. And Bits AI Investigation is already speeding up incident resolution and reducing expensive escalations. This customer will expand to 19 Datadog products.

Next, we signed a multiyear, over $30 million TCV deal with one of the world's largest online media companies. This customer chose to standardize on Datadog across its business, displacing four commercial and internal tools. Datadog also proved value beyond core observability with Product Analytics, CI Visibility, Data Observability and Cloud Cost Management.

This deal includes our largest win to date for Bring Your Own Cloud, displacing their legacy commercial logging tool at a petabyte scale. And finally, we signed a nine-figure renewal with a leading AI company. This longtime, very large customer uses 17 Datadog products to enable unified visibility on production workloads at a very large scale, albeit with a usage reduction starting in Q3, which we considered in our guidance and which David will speak to.

Before I turn it over to David for a financial review, let me offer 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 for our business. But we now have an additional growth driver with AI as we help our customers deliver value with this transformative new technology.

We are tremendously excited about our opportunities in AI. To summarize where we are and where we're going:

First, AI is a tailwind for Datadog today, as cloud consumption grows and drives more usage of our platform. As of Q2, over 750 AI customers use Datadog to monitor and improve their tech stacks. When we look at the largest companies driving AI, all 10 of the top 10 AI leaders are Datadog customers.

Beyond AI natives, we see AI activity growing across our broader customer base. We're also seeing signs of rapid growth in agentic activity, with the number of MCP tool calls quadrupling again quarter-over-quarter and growing more than 22x when compared to Q4 2025.

Second, we are delivering AI for Datadog, to deliver more value and greater platform capability to our customers. This includes our Bits AI products: Chat, Investigation, Detection, Code, Testing, Release, and many, many others.

Third, next-gen AI introduces new complexity and observability challenges. We are addressing this with what we call Datadog for AI, to observe and secure the AI stack from end to end. This includes GPU Monitoring, Agent Observability, Agent Console, Data Observability, AI Guard, and many other products.

Finally, our AI research team and our large volume of rich data used in critical workflows enable us to conduct groundbreaking research. We have shown some of our work already with the second version of our time-series model, TOTO, in May. TOTO version 2 was exciting for two reasons.

First, we've shown it to be state-of-the-art on key benchmarks. But more importantly, we've demonstrated for the first time true scalability for time-series models, allowing us to target the same improvement path language models have followed since 2020.

So now, beyond TOTO, we are working on larger and more ambitious dedicated models, post-training models to power Bits AI, and bringing other modalities beyond time-series data into world models that we think can lead to a step change in capabilities for our customers. And we plan to accelerate these research efforts with the acquisition of Adaptive ML, which we closed in June.

Because of all of that, now more than ever, we feel ideally positioned to help customers of every size and every industry, as well as all types of users, whether humans or AI agents, so they can transform, innovate and drive value through AI and cloud adoption.

And with that, I will turn it over to our CFO, David.

David M. Obstler
Chief Financial Officer

Thanks, Olivier. Our Q2 revenue was $1.12 billion, up 36% year-over-year. Within that, our 11% quarter-over-quarter revenue growth is the highest since Q2 2022, and our quarter-over-quarter revenue added of $115 million is a record by a significant margin. We continue to see robust usage growth from existing customers, as well as a strong ramp in our new customers.

Revenue growth accelerated with our broad base of customers, excluding AI customers, to the high-20s year-over-year, up from the mid-20s% last quarter and 18% in the year-ago quarter. We saw robust growth across our customer base with broad-based strength across customer size, spending bands and industries.

Meanwhile, our AI customers continue to grow rapidly and diversify in the quarter. This 750-strong customer group includes a broad range of AI startups, as it has in the past, but now also includes hyperscalers using Datadog for in-house AI labs. In Q2, this includes 31 customers spending more than $1 million annually, of which eight spent more than $10 million annually.

We also achieved strong new logo dollar bookings, with particular strength in enterprise, where new logo annualized bookings more than doubled from the year ago. And we are seeing new logos ramping faster and contributing more to revenue growth. The portion of our year-over-year revenue growth that relates to new customers was about 30% in Q2, up from about 25% in Q1.

Geographically, we're performing well in all regions, with growth acceleration across the regions. We see particular strength in the Americas, as much of the AI activity is occurring in the US, and as we execute strongly in LatAm.

Regarding retention metrics, our trailing-twelve-month net revenue retention percentage was in the low-120s, similar to last quarter, and churn remains low, with gross revenue retention in the mid-to-high 90s. We believe this metric highlights the mission-critical nature of our platform for our customers.

Now, moving on to our financial results. Billings were $1.18 billion, up 38% year-over-year. Remaining performance obligations, or RPO, was $3.47 billion, up 43% year-over-year. Current RPO grew about 40% year-over-year, and RPO duration increased year-over-year. As we previously mentioned, we continue to believe revenue is a better indicator of our business trends than billings and RPO.

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.

Our Q2 gross profit was $892 million for a gross margin of 79.6%. This compares to a gross margin of 80.2% last quarter and 80.9% in the year-ago quarter. As we've discussed in the past, our gross margin varies from quarter to quarter with investments into innovations for our customers, offset by efficiency efforts. There's no change in our expectations for gross margin, which has been in the 80% plus or minus range historically.

Q2 opex grew 26% year-over-year, versus 31% last quarter and 36% in the year-ago quarter. We held our DASH user conference in June, and as expected, the event cost about $15 million. Q2 operating income was $257 million for a 23% operating margin, compared to 22% last quarter and 20% in the year-ago quarter.

Turning to our balance sheet and cash flow statements, we ended the quarter with $5.0 billion in cash, cash equivalents and marketable securities. Cash flow from operations was $316 million in the quarter. After taking into consideration capital expenditures and capitalized software, free cash flow was $279 million for a free cash flow margin of 25%.

Now for our outlook for the third quarter and the fiscal year 2026. Our guidance philosophy overall remains unchanged. As a reminder, we base our guidance on trends observed in recent months and apply conservatism on these growth trends.

Regarding our largest customer, we have seen a usage reduction, which is incorporated in our Q3 and full-year 2026 guidance. As Olivier noted, this customer has recently renewed with us.

For the third quarter, we expect our revenue to be in the range of $1.135 billion to $1.145 billion, which represents 28% to 29% year-over-year growth. Non-GAAP operating income is expected to be in the range of $260 million to $270 million, which implies an operating margin of 23% to 24%. And non-GAAP net income per share is expected to be in the $0.63 to $0.65 per share range, based on approximately 378 million weighted average diluted shares outstanding.

For the full fiscal year 2026, we expect revenue to be in the range of $4.45 billion to $4.47 billion, which represents 30% year-over-year growth. Non-GAAP operating income is expected to be in the range of $1.01 billion to $1.03 billion, which implies an operating margin of 23%. And non-GAAP net income per share is expected to be in the range of $2.50 to $2.54 per share, based on approximately 376 million weighted average diluted shares outstanding.

And for some additional notes on guidance, we expect net interest and other income for the fiscal year 2026 to be approximately $180 million. We expect cash taxes in 2026 to be about $30 million to $40 million. We continue to apply a 21% non-GAAP tax rate for 2026 and going forward. And finally, we expect CapEx and capitalized software together to be in the 4% to 5% of revenue range in fiscal 2026.

Now, finally, to summarize, we are pleased with our execution in Q2. Our investments in R&D and go-to-market are yielding positive results, and they position us well for continued execution. I want to thank all Datadogs worldwide for their efforts.

And with that, we'll open the call for questions. Operator, let's begin the Q&A.

Question and Answer

Operator
Thank you. At this time, we will conduct the question-and-answer session. Our first question comes from the line of Sanjit Singh with Morgan Stanley. Your line is now open.

Sanjit K. Singh
Analyst, Morgan Stanley & Co. LLC

Thank you for taking the questions and congrats on the acceleration in revenue growth again this quarter. David, thank you for giving us the color on some of the guidance assumptions, particularly headed into Q3, with respect to the largest customer. I was wondering if you could share any additional details in terms of the new contract, was it of similar duration?

And in terms of the lower usage, is that a function of the customer getting a lower unit price because of making a new commitment? Or was there some churn or down-sell that we need to think through, not only for Q3, but for the balance of the year?

Olivier Pomel
Co-Founder, CEO & Director

Yeah. So maybe I'll take this one. I think, overall, as usual, we don't want to comment too much on any specific customer because we also don't really control what's happening with any specific customer. We wanted to be transparent about this on the call, because we did see a reduction in usage and we took the liberty to fully de-risk the guidance for the rest of the year with respect to that customer.

And again, the reason for that is we don't control what's happening to a specific customer, but we do have a great amount of control on what's happening to everything else in the business. And the business is booming. And we don't want that to overshadow basically the acceleration we see pretty much everywhere else in the business.

So as we mentioned on the call, we renewed the customer, it's a longtime customer who uses many of our products. But there's not a lot more we can share.

David M. Obstler
Chief Financial Officer

Yeah. I think...

Sanjit K. Singh
Analyst, Morgan Stanley & Co. LLC

Okay.

David M. Obstler
Chief Financial Officer

...just to get specific on the guidance. We, last quarter and previous quarters, said that we essentially have a level of commit and we can de-risk our guidance by using that. And then, as you know, in most of our large customers, we have variability relating to the commit, so we take that into consideration.

Olivier Pomel
Co-Founder, CEO & Director

Yeah. The last thing I'll say, because I know it's on people's minds, is if you back out our largest customer from our growth, you get pretty much the same growth rate, as the rest of the business has been accelerating very steadily. Actually we've seen, I think, now five quarters of continuous acceleration from the rest of the business. And we feel very good about what we see in the market.

Sanjit K. Singh
Analyst, Morgan Stanley & Co. LLC

Yeah. No, I appreciate the thought. Let's talk about maybe the rest of the business. What we've seen over the past couple of years, sort of AI-native leading the charge. It sounds like the enterprises are getting on board with their AI initiatives. And so just in terms of the enterprise AI, app dev cycle, what does that look like for Datadog over the last couple of quarters?

Olivier Pomel
Co-Founder, CEO & Director

Well, we do see broad adoption, and we see it in two ways. One is we see it manifest itself in just more transformation, more cloud adoption, more workloads, more modernization from customers. And that's what drives the majority of the known AI customer acceleration, so we mentioned also we've seen continuous acceleration from customers that existed before AI and that are not majority AI businesses. And that's been pretty remarkable. Like the acceleration that we gave the numbers on the call, but the acceleration since last year has been constant and very significant. And it keeps happening as far as we can tell. So it's a very positive trend there. That's the first thing we see.

The second thing we see is a very rapid increase in the usage of all of our AI-first services, so that would be the products that measure agents and LLMs. We see an explosion of traffic in terms of the LLM and tool calls we're getting. That would be the amount of calls we're getting to our MCP endpoints. So we see that explode completely over the past two quarters.

Sanjit K. Singh
Analyst, Morgan Stanley & Co. LLC

Appreciate the thoughts, Oli. Thanks.

Operator
Thank you. Our next question comes from the line of Raimo Lenschow with Barclays. Your line is now open.

Raimo Lenschow
Analyst, Barclays Capital, Inc.

Perfect. Thank you. Could I stay on that AI theme, please? At the moment, like if you think about the large customers, there's a lot of model training, et cetera. But if we broaden it out, inference is really becoming the bigger part. Can you talk a little bit about, like how much more observability is needed? And I'm thinking there if I do inference, I need to think about vector databases, I need to think about guardrails. All of these agents are going to be in containers that need to be monitored, et cetera.

So what do you see in real life at the moment in terms of if some people do more inference, how much more observability gets triggered by inference, is that kind of an opportunity that we should probably pay more attention to than that one renewal? And I had one follow-up.

Olivier Pomel
Co-Founder, CEO & Director

There's opportunity at every layer of the stack in inference. So we do think at the end of the day, inference will be the dominant workload, that's, anytime you train, you probably will want to infer more than you train as a rule of thumb. We see opportunity at the low level, when it comes to the infrastructure, the GPUs and the consumption you have there.

There's opportunities at the very top end when you measure what the agents are doing and whether you're getting the right outcomes or whether you're getting the right alignment, and there's opportunities at every layer in between, just looking at the LLM itself, just looking at the tool calls and the applications that are being called by the agents, like everything is an opportunity in there.

We see growing adoption from the products we already have there. We mentioned our GPU Monitoring product is actually getting quite a bit of usage in a number of neolabs and very AI-first types of customers. We're also seeing an explosion of volume in our agent monitoring product. And so we're well positioned there, but we think this market is going to change quite a bit and the preoccupations of customers will also change over time.

So for example, last year, our customers were mostly trying to validate correctness and validate that they were getting some form of outcome that it could then scale up. I would say three to six months ago, the focus has moved quite a bit towards cost.

Now customers were spending a lot on AI and they were wondering how to optimize cost. And I think we'll see variations in the concerns over time, as customers get further into the adoption and new problems emerge for them.

David M. Obstler
Chief Financial Officer

I just want to add that when you look at what we described as some of our deals in the quarter and you look down our scripts and you'll see that a number of them have the AI products included and so that is indication that those large enterprises are using the platform and buying the AI products as well.

Raimo Lenschow
Analyst, Barclays Capital, Inc.

Okay, perfect. Thank you. And then, David, one for you, it's like, obviously, you're always in a tough position if you have to guide and there's these large contracts. How did you do it historically? So did you always put in the base level and then what happened, happens? Or has that approach changed? I don't envy you on having to do this.

David M. Obstler
Chief Financial Officer

No. We essentially use, as we've talked about over the years, we kind of use the inputs of what we see. And what we've said, I think, in the last quarter or two, is that we have certain base levels. As you know, we have a commitment and a usage model, and we've factored that into providing our guidance.

So as we said in the prepared remarks, our methodology for guidance hasn't changed. We've always used those inputs and looked at the commitment and the usage in doing that.

Olivier Pomel
Co-Founder, CEO & Director

Yeah. The only thing I'll say is, in this case, we did choose to fully de-risk our largest customer. And the reason for that is we don't want that to be an overhang on what is otherwise a business that is accelerating and performing extremely well. So we have the same overall conservatism as we always do when we look at our numbers, but in this case we also [indiscernible]

David M. Obstler
Chief Financial Officer

Yeah.

Olivier Pomel

Co-Founder, CEO & Director

...this one a little differently.

Raimo Lenschow
Analyst, Barclays Capital, Inc.

Okay, perfect. That's very clear. Thank you.

David M. Obstler
Chief Financial Officer

Thank you.

Operator
Thank you. Our next question comes from the line of Gabriela Borges with GS. Your line is now open.

Gabriela Borges
Analyst, Goldman Sachs & Co. LLC

Hey, good morning. Thank you. I wanted to ask you both about one of our observations at DASH, which is the engineers love the pace of innovation. They talk very positively about the product. The CFOs love to complain a little bit about their Datadog bills.

So my question for you is talk to us a little bit about how those CFO level conversations are evolving. Clearly the ROI is there, but maybe give us a little more on where the budget is coming from. And something like Infinite Cardinality, is that now part of the conversation with CFOs in solving some of those very particular cardinality cost questions? Thank you.

Olivier Pomel
Co-Founder, CEO & Director

I mean, look, at a high level, there are only two reasons people buy software. It makes them more money or it saves them money. And anytime we sell, anytime we go out for a renewal, we go out for an upsell, or we land a new customer, that's because we do one of those two things for them. And we always have to make that case. So I wouldn't say that's any different from what we've seen before.

What we do for our customers today, especially as they keep adopting AI, is help them save a lot of the money they would spend on building, running operations, or running AI agents. When we have concerns with customers, that's the one thing they kept mentioning is, hey, how can you help me rein in my costs? This is growing very fast and I don't have any control on it. And I don't know whether I'm reaching the right outcomes with that. And so that's one of the reasons we've invested in all those products we've mentioned earlier. And also we're seeing some of the great returns on those products already.

In terms of Infinite Cardinality, that's, I would say, been one of the longest standing sources of frustration for customers, when sometimes they send more data or they send more fine grained tags with their data. And they get some unpredictability on the bills because of that, because it increases the cardinality of the data we're getting. And we've solved that from a technical perspective and from a commercial perspective by packaging our metrics a little bit differently.

And we think it's particularly important and relevant as customers are building more applications with AI. And as they want to send basically more tags, more information and more, ask more complex questions and get more fine grained answers to those questions. And so that fits well within their plans, basically. So we've got great feedback on that so far. But it's still early. Sometimes we get it right, sometimes we get it slightly wrong. And when we get it slightly wrong, we fix it. That's not different from what we've done in the past.

Gabriela Borges

Analyst, Goldman Sachs & Co. LLC

That all makes sense. Thank you for the detail.

Operator
Thank you. Our next question comes from the line of Mike Cikos with Needham. Your line is now open.

Mike Cikos
Analyst, Needham & Co. LLC

Hey, thanks for taking the questions, guys. I wanted to come back to the significant size of the lands that you had this quarter, and it's great to see the sustained traction, especially with those AI labs. But if I'm thinking about the two seven-figure AI labs that you landed this quarter, and then going to David's commentary around winning some of these in-house AI labs with the hyperscalers, are those one in the same here, or are those two separate customer sets we're talking about?

Olivier Pomel
Co-Founder, CEO & Director

These are different customers. The ones we mentioned on the new lands are neolabs. So these are companies that didn't exist a few years ago. And what's interesting about them, on the use case there, is that very often we land customers when they go into production and they release products and they start serving their customers.

In this case, these are customers who are getting, as they are training models, and they're using us to observe and improve and optimize the training of the models. And so that's an exciting new area. That was not really a business area for us a couple of years ago. And we've seen a number of new proof points around that.

In addition to that, and we've mentioned in previous calls, we've also landed the AI lab or super intelligence labs of a number of hyperscalers. And I would say the workloads are similar in that it's largely training of the models, but the customers are a bit different. These are very large companies that, in that case, previously had a lot of homegrown technology to observe and run workloads.

Mike Cikos
Analyst, Needham & Co. LLC

Excellent. And for a follow-up, I know you had cited the new logos ramping more strongly than what we've seen historically. And correct me if I'm wrong, but I feel like that's a newer phenomenon that you guys are calling out this quarter. When I think about those new logos ramping, is that a function of pull-through, where maybe some of these AI capabilities are pulling through the broader platform? Or is it vice versa? Anything you can do to help us think through what is creating that catalyst, if you will, when the new logos are contributing to the model? Thank you.

David M. Obstler
Chief Financial Officer

It's been happening and building up. So the number that we have in our Qs, which is the percent of growth from customers that we didn't have a year ago, that number, we said, has gone from 25% to 30%. So this has been building and we wanted to point that out because of that disclosure, indicating that the customers that we're landing, it's not only the new logos, but it's also the growth of the new logos that we've added over the last couple of years, last year, sorry. So it's a compounding of that.

Mike Cikos
Analyst, Needham & Co. LLC

Excellent. Thank you.

David M. Obstler
Chief Financial Officer

Thank you.

Operator
Thank you. Our next question comes from the line of Alex Zukin with Wolfe Research LLC. Your line is now open.

Alex Zukin
Analyst, Wolfe Research LLC

Hey, guys. Thanks for taking the question. Oli, maybe first for you, just on the, a lot of headlines around security over the course of the last few weeks, particularly AI breaking containment. And it occurs to me that with your positioning in observability and security, increasingly the notion of a guardian model and development around that could meaningfully increase kind of your ambit on what you can do and achieve for clients, both AI natives and legacy.

Can you maybe talk to what, the increasing opportunity around this crossover in this AI age and what that means for Datadog? And then I've got a quick follow-up for David.

Olivier Pomel
Co-Founder, CEO & Director

I mean, look, there's a complete switch in the way the security products need to work. So you can't wait basically for putting humans in the loop. You can't have the typical path when you have 12 or 15 different products that are going to aggregate signal, then you put that signal into a system to aggregate to prioritize them for humans, then humans will review them when they can. Like you need to integrate everything a lot more.

You need to operate a lot closer to the application and to the infrastructure, and you need to have AI agents solve the issues first. So it's a complete rebuild for most of the industry. And I think it plays into our approach, which is to have an integrated platform and have all of the different data streams come directly from observability straight into the security agents and have all that be integrated from end-to-end.

So obviously, this is a field that's moving very fast. We see new classes of issues pretty much every week at this point. We are quite busy building that up, but we think this plays into our strength and into where we already are and where we're building our security products.

Alex Zukin
Analyst, Wolfe Research LLC

Perfect. And then, David, maybe just for you. On the largest customer renewal, is there anything you can tell us around maybe just any changes around the duration or anything that makes this new contract maybe a little stickier in terms of the discounted rate card, the amount of products that they're able to kind of use for better value, anything that increases the conviction level around stickiness.

David M. Obstler
Chief Financial Officer

I won't comment on this other than to say that most of our enterprise customers, as we've talked about for a long time, have annual plus, and then the pricing is generally volume-based pricing. So I would say overall, our customers transact with us in that way. And then we have that level of commitment.

And then, as we've talked about over a lot of years, then there's usage and then we transact. So it's similar to what we have with most of our larger enterprise customers. Oli, anything you want to add there?

Olivier Pomel

Co-Founder, CEO & Director

No, I think there's a lot of continuity in that renewal. I think that's what you can, it's one way to put it.

Alex Zukin
Analyst, Wolfe Research LLC

Perfect. Thank you, guys.

Operator
Thank you. Our next question comes from the line of Eric Heath with KeyBanc Capital Markets. Your line is now open.

Ritu Kashyap
Analyst, KeyBanc Capital Markets, Inc.

Hi. This is Ritu Kashyap on for Eric Heath. I would love to hear more color on your 3Q guide specifically. It seems like it's a little below your sequential levels of how you've guided your previous 3Q. So we'd love to just hear more about what trends you're seeing going into 3Q, and maybe what some of the assumptions of the guide are.

David M. Obstler
Chief Financial Officer

Yeah, I think it's similar to the methodology we take, what we see and provide some conservatism. I think we had mentioned in the script that, while we've renewed our largest customer, we've seen usage declines relative to the previous quarter. So that's all taken into consideration in trying to develop a guidance that is consistent with the methodology of conservatism that we've used as a public company.

Ritu Kashyap
Analyst, KeyBanc Capital Markets, Inc.

Got you. And if I could just ask one more for Oli, I'd love to just get your thoughts on the impact of diversification of AI model usage in your customers and what you're seeing there?

Olivier Pomel
Co-Founder, CEO & Director

Well, we think it's great, there's a lot more options for customers to choose from in general. That opens up a lot of doors and opportunities for them. And it also creates a lot of complexity, and we're here to help deal with that complexity. So, for us, these are great opportunities.

And by the way, we've had that thesis since the early days of AI, that we would not just end up with one or two big AI companies and everybody using them the same way. We didn't just end up with one or two big cloud companies and everybody just using software from them. Like, the ecosystems are very, very, very rich. There are lots of providers, there are very large providers, there are smaller providers, and everything in between. And there are many compositions of those different systems that are used by any given customer. And so we think the same is going to happen in AI.

We think also that the multiplication of models and open source models in particular opens the door to customers doing a lot more training on their own. And so that's a new market for us. We see some signs that we have a very good role to play there. And so we're building towards that as well.

So overall, I mean, it's very positive for everyone.

Ritu Kashyap
Analyst, KeyBanc Capital Markets, Inc.

Got it. Thank you.

Operator
Thank you. Our next question comes from the line of Koji Ikeda with Bank of America. Your line is now open.

Koji Ikeda
Analyst, BofA Securities, Inc.

Yeah. Hey, guys, thanks so much for taking my question. Just one for me here. I wanted to ask on Bits AI, all the commentary that you guys are saying on Bits AI and all the work that we've been doing intra-quarter sounds like Bits AI is really taking off for you guys.

And so just thinking that Bits AI is going to be increasingly automating activities that historically has created observability workflows, I'm curious and really wonder, how do you ensure that greater automation that might be driven by Bits AI doesn't eventually reduce the volume of activity that traditionally drove Datadog consumption? Thank you.

Olivier Pomel
Co-Founder, CEO & Director

Well, look, if we provide more value, we'll get more. As I was saying earlier in the call, like, we sell more software by helping customers make more money or save money, or both. And I think if we can automate more and let them do more, we'll provide more value. That's as simple as that.

I think the future of observability is not just observing, it's fixing. It's not waking up people in the middle of the night because something broke, but fixing it for them. It's not letting people do damage control on a security incident because an attacker is in, it's preventing the attacker from getting in to start with, by auto-remediating issues. And we're very, very busy building all of that. And we're super confident that this will yield great business outcomes for us in the end. And that's what we see from customers in the market. Like, when they use Bits AI, they use more of our product, they deploy more of it, they create more dashboards and alerts and everything else. They have more users inside of our product. Like, it's not a zero-sum game.

Koji Ikeda
Analyst, BofA Securities, Inc.

Thank you.

Operator
Thank you. Our next question comes from the line of Samik Chatterjee with JPMorgan. Your line is now open.

Samik Chatterjee
Analyst, JPMorgan Securities LLC

Great. Thanks. And thanks for taking my question. Maybe just on the non-AI part and the acceleration that you're seeing related to the non-AI part of the business, just wanted to sort of get your thoughts on the sustainability, and whether this acceleration that you're seeing is driven by some of the new customer logos that you're pointing out, or more usage going up.

And as CFOs get more sort of cautious around their budgets, do you see more sensitivity around non-AI eventually relative to some of the AI products and how they're doing at this point? And I have a quick follow-up. Thank you.

Olivier Pomel
Co-Founder, CEO & Director

I mean, from what we can tell, it's very broad-based. And it's largely driven by existing customers, because that's the majority, like, when you think of what it takes to move that number, that's basically the majority of our business. Like, we're not just going to move that with a few newer customers there, it's largely driven by the existing customers. And it's driven by both increases in volume, because they're moving more workloads to the cloud, and adoption of our newer products as they consolidate onto us.

We think it's sustainable for one thing, if you compare it to what we have seen in the heady days of 2021, or the growth rates are accelerating, but they're still far below what we were seeing at that time. And so we don't create the same issue of customers having to digest very large increases multiple years in a row. I think in this case, we're very well within the realms of sustainability.

And as has been a theme in this call, remember that when customers adopt and they consolidate, they have an eye towards the financial side of the equation basically, how much money are they going to make or save by doing that at the end. And we are very good at helping customers understand that and making that case and helping them save money at the end of the day. So we feel good about that.

David M. Obstler
Chief Financial Officer

And I want to just add one thing, and we talked about this last quarter, that some of this has to do with the investments that we're making in our platform and our product, but it also has to do with the investments that we're making in our go-to-market. We've successfully expanded quota capacity, the geography of it, and essentially, that's, as we talked about last quarter, providing returns. So that's also being a growth driver in our non-AI or enterprise type business.

Olivier Pomel
Co-Founder, CEO & Director

That's right. And you see it also contributing to growth there. So we keep investing in R&D, obviously, because we're shipping more products that are successfully being adopted and consolidated into by a large number of existing customers. But we also are adding to our go-to-market teams. We're still not at the scale we want to be in terms of getting into all of the customers worldwide in all of the segments that are relevant to us. So we're investing as we see the return on those investments.

Samik Chatterjee
Analyst, JPMorgan Securities LLC

Got it. And for my quick follow-up here, you talked about the FedRAMP High certification last quarter. Just curious if there's anything to sort of update us on the pipeline, and if there's any momentum on that front on the pipeline yet? Thank you.

Olivier Pomel
Co-Founder, CEO & Director

Yeah. We're investing quite a bit in the build-up of our federal and government sale in general. And we see pipeline there. In general, these are not deals that happen overnight. But this is a very large market. And we see great traction there. And we're investing to take full advantage of it.

A lot of that was a build-up to get to the right level of certification, so we can deliver SaaS to various levels of government. And we've done quite a bit there. There's actually even more we're planning to do there. But we're happy with the results so far.

Samik Chatterjee
Analyst, JPMorgan Securities LLC

Thank you.

Operator
Thank you. Our next question comes from the line of Howard Ma with Guggenheim Securities. Your line is now open.

Howard Ma
Analyst, Guggenheim Securities LLC

Great. Thank you. And congrats on the strong quarter and full year guidance raised. I have two questions. I'll just ask them together. The first is on Bits AI. I'm curious how adoption and contribution compares to previous major feature expansions in the past.

And then my other question is the $30 million TCV deal with the, I think you guys said it's one of the largest online media companies. I'm assuming this company did mostly DIY before. So if you could share some light on the decision making process and if they're using multiple Datadog products and why now, that'd be really helpful. Thank you.

Olivier Pomel
Co-Founder, CEO & Director

Yeah. I'm sorry, I missed some part of your second question.

Yuka Broderick
SVP of Strategic Finance & Investor Relations

It was, are they taking multiple products? I think, right, Howard?

Howard Ma
Analyst, Guggenheim Securities LLC

Yeah, media. Are they...

Olivier Pomel
Co-Founder, CEO & Director

Yeah.

Howard Ma
Analyst, Guggenheim Securities LLC

Yeah. The nature of the sale...

Olivier Pomel
Co-Founder, CEO & Director

Nature of the sale...

Howard Ma
Analyst, Guggenheim Securities LLC

Yeah.

Olivier Pomel
Co-Founder, CEO & Director

Why now?

Howard Ma
Analyst, Guggenheim Securities LLC

Yeah. That client, yeah.

Olivier Pomel
Co-Founder, CEO & Director

Yeah. So first on Bits AI, so yes, one thing that happened is Bits AI used to be fairly specific. It used to be dedicated to alerts. Bits AI would pick up an alert and would run an investigation for you.

Now the surface of contact is a lot wider with a customer. So Bits AI, you can access it through chat. You can, or you can, of course, still do the investigations, and we've done quite a bit more there.

You can have Bits AI manage your monitoring and manage your detection for you. You can have it code for you, you can have it generate managed tests. There's all sorts of different use cases that we've built into it that broadened the surface of contact. And we see a lot of adoption across all of those different areas.

We're also changing the way we package it. So we have a new model with AI credits that we're rolling out, just because the surface of contact is so much wider now than the specific feature. So there's quite a bit that is going on there. The explosion of activity that I mentioned earlier about other AI services is happening also in Bits AI. So that's something we're looking forward to.

That's on that. On the second one, on the products that are being adopted in the sale, I mean, look, we typically land with two or more products. The balance we try to strike there is always to land enough of the platform without slowing down the deals too much, because the more you try to do at once, the more stakeholders you get in, the longer it takes. And so we found that two products in general is a good land. And then we can expand from there.

On the calls, we tend to mention a lot of consolidation deals, because they tend to be the larger ones. If you land with 12 products, you're going to be larger than if you land with two, in general. That's not the majority of the deals. The consolidation typically happens later than when we land, but this makes for very interesting examples of what our customers are doing when they're consolidated on us all at once.

Operator
Thank you.

Howard Ma
Analyst, Guggenheim Securities LLC

Thank you.

Operator
Our next question comes from the line of Andrew Sherman with TD Cowen. Your line is now open.

Andrew Sherman
Analyst, TD Securities (USA) LLC

Great. Thank you. And congrats on the core growth acceleration. Oli, CPUs have had a renaissance lately driven by agentic AI, would be great to hear your thoughts on this topic. If it can be an incremental growth driver for your infrastructure monitoring, have you seen any evidence of this yet? That's it for me. Thanks.

Olivier Pomel
Co-Founder, CEO & Director

Look, we do see an acceleration of consumption of our infrastructure products in general. And also that's, at a high level, we do see that across the customer base. I don't know that we see specifically the CPUs that get attached to GPUs in the new build out. I think a lot of it has more to do with the fact that the AI agents are largely spending a good amount of their time, like sometimes the majority of their time, calling tools, and tools are just applications that already existed. And those applications typically run on CPUs. And so we see quite a bit of that.

Operator
Thank you. Our next question comes from the line of Brad Reback with Stifel. Your line is now open.

Brad Reback
Analyst, Stifel, Nicolaus & Co., Inc.

Great. Thanks very much. Oli, given your commentary around how strong the core is and that your largest customer was not additive to growth here in 2Q, should we assume that if we ex-out the sequential down-tick in that customer, that the core guide would have been probably 300 basis points or 400 basis points higher?

Olivier Pomel
Co-Founder, CEO & Director

Well, I can't speculate, but what I will say is, look, the business overall is growing at the same rate if you exclude that customer, as I said. And the business has been accelerating overall. That's why we feel good. Like, when we look at whether we're getting the right returns and the right outcomes for our investments in R&D or investments in go-to-market, when you look at our pipelines and all of the signs we have about the business, we feel great about the business. So it's a good time to be in business.

David M. Obstler
Chief Financial Officer

Yeah. I think we commented in the remarks that the non-AI has accelerated and the AI, excluding the largest customer, continues. So I think we gave those trends in describing the business.

Olivier Pomel
Co-Founder, CEO & Director

Customers are growing a lot faster than non-AI.

David M. Obstler
Chief Financial Officer

Yeah. And that's...

Brad Reback
Analyst, Stifel, Nicolaus & Co., Inc.

Thanks. Great.

David M. Obstler
Chief Financial Officer

...the AI is growing. Yeah. Exactly.

Brad Reback
Analyst, Stifel, Nicolaus & Co., Inc.

Perfect. Thank you, guys.

Operator
Thank you. Our next question comes from the line of Ittai Kidron with Oppenheimer & Co. Your line is now open.

Ittai Kidron
Analyst, Oppenheimer & Co., Inc.

Thanks, guys. And congrats guys on a great quarter. I wanted to ask about new customer additions. This probably was the weakest quarter I ever remember for you guys, especially in a quarter where you had

DASH, where historically DASH has been an accelerant of new customer additions. Any color there would be great.

Olivier Pomel
Co-Founder, CEO & Director

Dave?

David M. Obstler
Chief Financial Officer

Yeah. I think essentially, it's very similar to what we talked about before. Our gross customer additions continue to be strong and on trend line. And that's the vast majority of our revenues. We have, at the very low end, the border between free and contract, and that has variability, very low effect on revenues. So if you, that accounts, as we talked about in many quarters, that accounts for the variability of the customer count. And it really has to do with something that has very little effect on revenues.

Olivier Pomel
Co-Founder, CEO & Director

Yeah. When you look at the customers above certain thresholds, like whether it's above $1 million, above $100,000, above $10,000, like, we know all of those are trending very well.

Ittai Kidron
Analyst, Oppenheimer & Co., Inc.

Very good. And then as a follow-up, Oli, for you, perhaps I want to follow up on the questions around Bits, which sounds super interesting. I guess the longer term, and as you try to push deeper, Oli, also into the security side of things, could this evolve into some broader AI SOC automation kind of platform? Is that a reasonable direction to think that this is where it's going to go?

Olivier Pomel
Co-Founder, CEO & Director

Well, there's definitely, we've taken moves towards that, right? So we initially, we built, so we built a SIEM first for that. Then we built the agent into the SIEM, so Bits AI Security Analyst. And now we've actually separated the agent from our SIEM, so customers can use it with other SIEMs. And we do that because the agent performs so well, and it's been such a differentiator when we pitch the SIEM, that we think we're limiting ourselves market-wise if we just go after customers that want to re-platform their SIEM, and it can have a much broader appeal as an AI SOC. So we are definitely taking moves towards that.

Ittai Kidron
Analyst, Oppenheimer & Co., Inc.

Very good, I appreciate it. Thank you.

Operator
Thank you. Our next question comes from the line of Andrew DeGasperi with BNP Paribas. Your line is now open.

Andrew DeGasperi
Analyst, BNP Paribas Securities Corp.

Thanks for fitting me in. I just wanted to ask a question on the non-AI natives, specifically in terms of the growth that you saw in the quarter. I was wondering, did you see rising demand for the AI monitoring tool, particularly with open source tools being deployed across enterprises?

Olivier Pomel
Co-Founder, CEO & Director

Sorry, I missed the center part of your question. The...

Andrew DeGasperi
Analyst, BNP Paribas Securities Corp.

In terms of the AI monitoring...

David M. Obstler
Chief Financial Officer

I think you're asking about within that, the AI, what we used to call AI monitoring, I think you're asking...

Yuka Broderick
SVP of Strategic Finance & Investor Relations

LLM.

David M. Obstler
Chief Financial Officer

LLM et cetera, the growth trend there.

Andrew DeGasperi
Analyst, BNP Paribas Securities Corp.

Yeah.

David M. Obstler
Chief Financial Officer

Yeah.

Olivier Pomel
Co-Founder, CEO & Director

And look, the volume, like, there used to be very little volume a year ago, it started growing quite a bit into the second half of last year. And now it's been very rapidly accelerating over the past couple of quarters. So we've seen an explosion basically of the volume we're getting there.

And we get more usage from different kinds of companies, so we definitely see that. We see it also across traditional companies and some more recent AI natives. So we see a little bit of both. I would say for that category, it's still super, super early, like we expect the products to change quite a bit. We expect the usage and maybe also the packaging to change over time quite a bit.

Andrew DeGasperi
Analyst, BNP Paribas Securities Corp.

Got it. Thank you.

Olivier Pomel
Co-Founder, CEO & Director

All right. So I think that was the last question. So, I want to thank all of you for attending the call today. I also want to again thank the teams, everyone at Datadog. I think everybody has been doing a fantastic job, both on the product side and go-to-market side. I know we have a lot more lined up for the end of the year on the product side. And I know also we have very large and very happy pipelines to tend to on the go-to-market side. So I hope to talk to you again next quarter. Thank you all.

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