

Datadog, Inc. NasdaqGS:DDOG

FQ1 2026 Earnings Call Transcripts

Thursday, May 7, 2026 8:00 AM ET

CONTENTS

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

SVP of Strategic Finance & Investor Relations

ANALYSTS

Fatima Aslam Boolani

Citigroup Inc., Research Division

Gabriela Borges

Goldman Sachs Group, Inc., Research Division

Gregg Steven Moskowitz

Mizuho Securities USA LLC, Research Division

Karl Emil Keirstead

UBS Investment Bank, Research Division

Mark Ronald Murphy

JPMorgan Chase & Co, Research Division

Patrick Edwin Ronald Colville

Scotiabank Global Banking and Markets, Research Division

Peter Weed

Bernstein Institutional Services LLC, 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 First Quarter 2026 Datadog Earnings Conference Call. [Operator Instructions] Please be advised that today's conference is being recorded. I would now like to turn the conference over to Yuka Broderick, Senior Vice President of Investor Relations. Please go ahead.

Yuka Broderick

Investor Relations

Thank you, Lisa. Good morning, and thank you all for joining us to review Datadog's first 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 second 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 and 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-K for the year ended December 31, 2025. Additional information will be made available in our upcoming Form 10-Q for the fiscal quarter ending March 31, 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 a very strong start to 2026. Let me begin with this quarter's business drivers. I'm very pleased to say that our teams executed very well and delivered revenue growth of 32% year-over-year, accelerating from 29% last quarter and 25% in the year ago quarter. We showed broad-based acceleration of revenue growth across cohorts, including both our AI and non-AI customers. Our AI-native customer cohort continues to grow and diversify rapidly both in the number of customers we serve and the scale of those customers. And this quarter included new land deals with 2 of the world's biggest AI research teams, helping them improve and optimize their training workflows. I'll talk more about that in a bit.

Even more impressive was the growth in our non-AI customers. non-AI customer revenue growth accelerated again this quarter to mid-20% year-over-year up from 23% last quarter and 19% in the year ago quarter. We think this is a sign of strong continued cloud migration, greater adoption of our products, and customers of all kinds accelerating their use of AI. Finally, 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 Q1 financial performance and key metrics. Revenue was \$1.01 billion, an increase of 32% year-over-year and above the high end of our guidance range. We ended Q1 with about 33,200 customers up from about 30,500 a year ago. We also ended with about 4,550 customers with an ARR of \$100,000 or more, up from about 3,770 a year ago. These customers generated about 90% of our ARR. And we generated free cash flow of \$289 million with a free cash flow margin of 29%.

Turning to product adoption. Our platform strategy continues to resonate in the market. For example, 56% of our customers now use four or more products, up from 51% a year ago. 35% of our customers use 6 or more products, up from 28% a year ago, and 20% of our customers use 8 or more products, up from 13% a year ago. So, we're landing more customers and delivering value across more products. And our business continues to grow. Our total ARR now exceeds \$4 billion, and our quarterly revenue exceeded \$1 billion for the first time. This is a big achievement for all of us at Datadog and is a product of years of investment in building and innovating for our customers. But we are still just getting started. Of our 26 products, 5 are over \$100 million in ARR and another 3 are between \$50million and \$100 million ARR. We're working hard to build and deliver further growth in those products. And this leaves 18 other products, which are earlier in their lifecycles. We believe each has the potential to grow to more than \$100 million over time.

Moving on to R&D. Our engineers, enabled with the latest AI coding tools, are building rapidly to help our customers confidently and securely deploy their applications. So let me speak to a few of our product launches this quarter. Let's start with AI. As a reminder, we're talking about our AI efforts in 2 buckets: AI for Datadog and Datadog for AI. So first, AI for Datadog. These are AI products and capabilities that make the Datadog platform better and more useful for our customers. In March, we launched our MCP server for general availability. With MCP Server, developers access live production data to debug their applications directly in their AI coding agent or IDE. We delivered Bits AI Security Agent, which autonomously triages Datadog Cloud SIEM signals, conducts in-depth investigations of potential threats, and delivers actionable recommendations. We've seen Bits AI Security Agent reduce investigations that could take hours to as little as 30 seconds. We also shipped Bits Assistant now in Preview, which helps customers search and act across Datadog using natural language prompts.

Moving on to Datadog for AI. This includes Datadog capabilities that deliver end-to-end Observability and security across the AI stack. We launched GPU monitoring, enabling teams to understand GPU fleet utilization, workload efficiency, thermal and power behavior and interconnect performance. This drives higher GPU ROI and operational reliability. Our customers continue to move forward with their AI activities, and we can see that in their usage of the Datadog platform. We now have over 6,500 customers sending data for 1 or more AI integrations. Though this is only 20% of total customers, they represent about 80% of our ARR. And our customers' usage of AI within the Datadog platform continues to grow rapidly. Bits AI SRE Agent investigations have more than doubled from December to March. The number of spans sent to our LLM Observability product nearly tripled quarter-over-quarter. The number of Datadog MCP Server tool calls, quadrupled quarter-over-quarter and the number of Bits Assistant messages increased by a factor of twelve in that period. While we are aggressively building with and for AI, we also continue to expand the Datadog platform to deliver against our customers' increasingly complex needs. To speak to a few of these efforts.

Last month, we launched experiments for general availability. Experiments work hand-in-hand with our feature flagging product and combine best-in-class statistical methods with real time observability guardrail among alternatives so companies can test for impact, choose among alternatives quickly and ship with confidence. In addition, our customers now benefit from APM recommendations by analyzing telemetry data from application performance monitoring, Real User Monitoring, profiler and database monitoring recommendations, APM automatically identified performance and reliability issues and most importantly, explain how to fix them. And we announced our plans to launch our next datacenter in the U.K. We see a large opportunity to serve our British customers as cloud adoption accelerates in regulated industries. Last but not least, we are pleased to have received FedRAMP High certification from the U.S. federal government. With this certification, we can now move forward with federal agency customers that require FedRAMP High to handle sensitive workloads. Meanwhile, we continue to expand our product offerings, go-to-market teams and channel partnerships for public sector customers, both in the U.S. and internationally. So, our teams were hard at work again. and we're looking forward to sharing many new products and future announcements at our DASH conference on June 9 and 10 in New York City.

Now let's move on to sales and marketing and highlight some of the deals we closed this quarter. First, we landed 2 large deals, a 7 figure and an 8-figure annualized deals with the AI research divisions at 2 of the world's largest technology companies. These organizations are building and training the most advanced AI models in the world. It is critical for them to reduce engineering friction and increase training

velocity. But fragmented internal and open-source tooling made it harder to identify and solve issues, and reduced engineering and research productivity. By using Datadog, both companies are accelerating their pace of innovation on their hyperscale AI training workloads. And this includes optimizing their workflows using GPU monitoring on large parallel GPU grids.

Next, we signed a 7-figure annualized expansion for an 8-figure annualized deal with a leading online recruiting platform. This customer is centralizing on Datadog to reduce complexity, drive developer velocity and improve efficiency. With this expansion, they will replace a standalone tool with Datadog LLM Observability to correlate LLM signals with APM and user experience data. This customer will grow to 16 Datadog products, including Datadog MCP Server.

Next, we signed a 7-figure annualized expansion for an 8-figure annualized deal with a Fortune 500 bank. With this expansion, this customer will migrate the remaining log data into Datadog, fully replacing their legacy log vendor. Most notably, our Flex logs give them granular control over costs while meeting strict compliance requirements. This customer uses 10 Datadog products, including Bits AI SRE Agent, to accelerate incident response with AI.

Next, we signed a 7-figure annualized expansion with a leading global hedge fund. This customer operates thousands of on-prem hosts and network devices. At that scale, their open-source monitoring stack had become operationally unsustainable, impacting portfolio managers and investment analysts. With this expansion, they will replace their entire on-prem Observability layer with Datadog infrastructure monitoring and network device monitoring and will have unified visibility across their cloud and on-prem environment. This customer will expand to 11 Datadog products.

Next, we landed a 6-figure annualized deal with a Fortune 500 insurance company. This company's fragmented observability stack led to long outages with incidents reported first by their customers instead of their tooling. By using Datadog and consolidating 3 legacy APM tools, they expect to move from reactive responses to proactive incident detection. They will adopt 10 Datadog products to start, including all 3 pillars and LLM Observability.

Next, we signed a 7-figure annualized expansion with one of the world's largest travel groups in APAC. This customer was using Datadog in one business unit, but in 2 others, they were juggling multiple tools and lacked actionable insights. By consolidating 6 legacy open source and cloud monitoring tools, the customers save money and improve platform resiliency and performance. This multiyear commitment positions Datadog as their strategic observability provider.

And finally, we landed a 6-figure annualized deal with a leading Latin American fintech company. This customer serves tens of millions of users across critical financial flows. Their rapid growth outpaced their fragmented front-end monitoring setup, and outages exposed them to financial, operational and reputational risks. By adopting our digital experience monitoring suite including RUM, Synthetics and product analytics, they now have full visibility over user activity with the cost controls they also previously lacked. This customer will start with 5 Datadog products.

And that's it for our wins. Congratulations again to our entire go-to-market organization for a great Q1. Before I turn it over to David for a financial review, I want to say a few words on our longer-term outlook. We are pleased with the way we started 2026 as we support our customers' inflection in AI usage and application development and as they lean into our AI innovations, including Bits AI SRE Agent, Bits AI Security Analyst, Bits Assistant, Datadog MCP Server, GPU monitoring and many more. 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 secular growth driver with AI as we help our customers deliver value with this transformative new technology. 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. This was a very strong quarter for Datadog. Our Q1 revenue was \$1.01 billion, up 32% year-over-year. Our 6% quarter-over-quarter revenue growth is the highest for Q1 and since 2022. And our \$53 million quarter-over-quarter revenue added is the highest ever for Q1. That included the strongest quarter of sequential usage growth from existing customers since the first quarter of 2022. We also delivered an all-time record for sequential ARR added in a quarter. ARR growth accelerated in each month of Q1, and we see a continuation of these healthy growth trends in April.

We also achieved strong new logo bookings. New logo annualized bookings set a new all-time record by a significant margin and more than doubled versus a year ago quarter. These included wins in observability and included some of our newer products like security, data observability and Flex logs. And our new logo average land size also set a record and more than doubled year-over-year as we continue to land larger deals. Revenue growth accelerated with our broad base of customers, excluding the AI natives to mid-20s percent year-over-year, up from 23% last quarter and 19% 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 native customer growth continues to significantly outpace the rest of the business. This group continues to diversify and grow, including 22 customers spending more than \$1 million annually, and five, spending more than \$10 million annually. This group includes the leading companies in foundational models, code gen tools and vertical-specific AI solutions.

Next, regarding our retention metrics. Our trailing 12-month net revenue retention percentage was in the low 120%, up from about 120 last quarter and our trailing 12-month gross retention percentage remains in the mid- to high 90s.

Now moving on to our financial results. Billings were \$1.03 billion, up 37% year-over-year and remaining performance obligations, or RPO, was \$3.48 billion, up 51% year-over-year, with current RPO growing in the mid-40s percent year-over-year. RPO duration increased year-over-year as the mix of multiyear deals increased in Q1. As a reminder, we continue to believe revenue is a better indicator of our business trends than billings and RPO given their variability.

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, Q1 gross profit was \$807 million, with a gross margin of 80.2%. This compares to a gross margin of 81.4% last quarter and 80.3% 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. Our Q1 OpEx grew 31% year-over-year versus 29% last quarter and 29% in the year ago quarter. As a reminder, we continue to grow our investments to pursue our long-term growth opportunities, and this OpEx growth is an indication of our execution on our hiring plans. Q1 operating income was \$223 million or a 22% operating margin compared to 24% last quarter, and 22% in the year ago quarter.

Turning to the balance sheet and cash flow statements. We ended the quarter with \$4.8 billion in cash, cash equivalents and marketable securities. Our cash flow from operations was \$335 million in the quarter. After taking into consideration capital expenditures and capitalized software, free cash flow was \$289 million and free cash flow margin was 29%.

And now for our outlook for the second quarter and for the fiscal year 2026. First, our guidance philosophy overall remains unchanged. As a reminder, we base our guidance on trends observed in recent months and apply conservatism on these growth trends. In addition, as with last quarter, we are applying a higher degree of conservatism to our largest customer. So, for the second quarter, we expect revenues to be in the range of \$1.07 billion to \$1.08 billion, which represents a 29% to 31% year-over-year growth. This guidance implies sequential revenue growth of \$64 million to \$74 million or 6% to 7%, due to the strong growth of revenue in Q1 and into April. Non-GAAP operating income is expected to be in the range of \$225 million to \$235 million, which implies an operating margin of 21% to 22%. As a reminder, in Q2, we will be holding our DASH user conference which we estimate to cost about \$15

million, which we have reflected in our operating income guidance. Non-GAAP net income per share is expected to be \$0.57 to \$0.59 per share based on approximately 369 million weighted average diluted shares outstanding.

And for fiscal 2026, we expect revenues to be in the range of \$4.3 billion to \$4.34 billion, which represents 25% to 27% year-over-year growth. Non-GAAP operating income is expected to be in the range of \$940 million to \$980 million, which implies an operating margin of 22% to 23%. And non-GAAP net income per share is expected to be in the range of \$2.36 to \$2.44 -- \$2.36 to \$2.44 per share based on approximately 372 million weighted average diluted shares outstanding. Finally, some additional notes on the guidance. We expect net interest and other income for fiscal 2026 to be approximately \$170 million. We expect cash taxes for 2026 to be approximately \$30 million to \$40 million. We continue to apply a 21% non-GAAP tax rate for 2026 and going forward. And we expect capital expenditures and capitalized software together to be 4% to 5% of revenue in fiscal 2026.

To summarize, we are very pleased with our execution in Q1. We are well positioned to help our existing and prospective customers with their cloud migration, digital transformation, and AI adoption journeys. And I want to thank Datadogs worldwide for their efforts. With that, we'll open the call for questions. Operator, let's begin the Q&A. Thanks.

Question and Answer

Operator

[Operator Instructions] Our first question today is coming from the line of Mark Murphy of JPMorgan.

Mark Ronald Murphy

JPMorgan Chase & Co, Research Division

Congratulations on an amazing performance. Olivier, is there any way to conceptualize the growth in the sheer raw volume of code being produced in the world today due to adoption of code generators such as Claude Code and Codex and cursor, because they seem to be developing the capability to take on full projects and some of the charts are showing these capabilities are just exponentially exploding upward in a straight line. I'm wondering how much of that code is going into production and therefore, driving activity for Datadog.

Olivier Pomel

Co-Founder, CEO & Director

Well, we definitely think and see that here's many more applications being created. There's going to be way more complexity in production. We see some of that happening already today. Some of those new applications are getting into production, they're finding users. We see some signs of that at every layer of our platform. We quoted a few stats on the increasing data volumes. We see AI products; that's definitely a reflection of that. So, we see an inflection point there in consumption from customers. We see a move to production that is very real, and we see that across both AI-native and non-AI companies.

Mark Ronald Murphy

JPMorgan Chase & Co, Research Division

Okay. And as just a quick, related follow-up. If we click down one layer, and I'm wondering how you might view the increasing heterogeneity of the environment at the silicon level, because the -- when you look across Amazon with Trinium and Graviton and Google with TVs, Microsoft has launched the myosilicon. It looks like that is starting to explode. Our understanding is that trying to monitor the mixed environment is a lot more difficult than if you just have a uniform fleet of Intel and AMD chips, and we keep hearing all the traditional monitoring tools, they really fail on the custom silicon and Datadog handles it well. The -- and then all this new telemetry, including high-bandwidth memory and that type of thing. Can you speak to whether that trend is giving you some tailwinds?

Olivier Pomel

Co-Founder, CEO & Director

Yes. I mean, look, broader market that's interesting here is if it's training, training used to be something only 2 or 3 companies were doing or maybe 4, 5 at a large scale. And it looks like training might democratize quite a bit more, and many companies will train models on a regular basis. So, it becomes more of a viable category for service providers -- selling provider like us basically. I think the heterogeneity of the silicon is definitely a trend that plays in our favor there. The more heterogeneous, the more you need someone else to make sense of everything for you and also stitch together with the non-GPU aspects and the rest of the infrastructure, the application, the users, and the developers like basically everything we are used for. When you think of who actually has heterogeneous environments today, that is still a very small number of companies. Google barely just started selling their TPUs to the outside. So I think it's still a small number of companies that are there, but we see a growing opportunity there. Interestingly, last year, when we reported earnings, we said we're mostly interested in inference workloads and training is not a real market for us yet. Now we actually see training becoming a market. We started landing customers that are actually hyperscalers that have a whole host of homegrown technologies and that are using us specifically in their super intelligence labs to help monitor their

workloads, accelerate the training runs, monitor the GPUs also. So, we see that as a point of validation that there's going to be a fit for us

Mark Ronald Murphy

JPMorgan Chase & Co, Research Division

That's amazing to think there's a whole need to mention, if you can move from inferencing into the training side. And I caught the reference in the prepared remarks of how you landed a couple of those very large labs. So congrats on everything.

Operator

And our next question will come from the line of Sanjit Singh of Morgan Stanley.

Sanjit Kumar Singh

Morgan Stanley, Research Division

I want to spin off with David on this guide to start the year is probably the best we've seen in several years, David, and laid out the underlying assumptions quite well. Just wanted to do a sanity check just on the sort of overall backdrop macro backup, we do have some geopolitical tensions and those types of things when we think about. Your Mid-East Base business and any impact from like in your e-commerce or retail business, where there may be some consumer discretionary impacts. I just want to get like how you're thinking about those parts of the business. And then I had a follow-up for Oli.

David M. Obstler

Chief Financial Officer

Yes. We had a very strong quarter across the board. We have a multi-industry multi-geography type of quarter, and SMB was very strong. And that -- the source of our guidance and our raises are at the core, that type of performance. We haven't seen particular effect in the consumer businesses or e-commerce businesses yet. We basically have a continuation of trends in those businesses, travels and things like that. that are very similar to the other industry. So, we haven't seen it yet. We obviously watch it and look at analytics, but we haven't seen it. In terms of our overall guidance, the trends that we have in organic, we discount across the board, and I think we mentioned our particular treatment of our largest customer.

Sanjit Kumar Singh

Morgan Stanley, Research Division

That's very clear. And then Olivier, for you. I think we -- when we talk to investors about the debate in this category longer term is just what does this what does the category look like when agents are doing the triaging investigating versus human engineers and human SREs. And so -- what is your sort of vision of that -- how that evolves for Datadog, both from a product standpoint and an experience standpoint from a UI perspective, but also like is there going to be immune modalities in terms of pricing when agents are consuming the Datadog platform to a higher degree than engineers do today?

Olivier Pomel

Co-Founder, CEO & Director

Yes. Look, I think one thing I'd say is it's hard to tell where we're going to be in 4 or 5 years. If you had told me 2 years ago that most engineers would go back to coating in the console. I wouldn't have believed you. And yet, that's one of the winning modalities today. Look, as far as we're concerned, we don't care whether most of the usage is humans, most of usages agents. Our business model lends itself to do well. We are usage-based; it doesn't really matter where the is coming from that perspective. The way we see trends right now is, we see both a stratospheric increase of agent usage. So, we have a ton of usage on our MCP server. We see customers spending to automate a lot with their own agents using our agent combination of those. But we also see an increase of usage of the web interface by humans. So right now, the two work hand-in-hand and we keep developing and pushing on those fronts.

Operator

Next question is coming from the line of Raimo Lenschow of Barclays.

Raimo Lenschow

Barclays Bank PLC, Research Division

One for Olivier, one for David. Olivier, if I listen to you in your prepared remarks, there's a lot of like consolidation that people try to do open-source tooling and then realize they kind of needed to come to you and come back. On the other hand, in the industry, we still have a lot of like noise around that level. How do you see it in real life. To me, it seems a little bit like optionability is just very hot. And then there's different categories where you use certain items -- certain vendors and some open source, can you speak what you see in real life there?

Olivier Pomel

Co-Founder, CEO & Director

I mean, in real life, most companies have open source in some capacity somewhere. When it comes to having a platform that unifies everything telecare everything does more of the problem solving for you, that's typically what customers use us. And the motion we see pretty much everywhere, these customers have 4 or 6, 7, 15, and 25 different things, and different pockets in the organization, and different business units, and it's a huge mass. And they come to us, they can unify all that. They get better results because all of the data is in one place, the workflows can be automated from end-to-end—you can get end-to-end visibility, you don't have blind spots. And also, they save money because they don't have all these pockets of inefficiency everywhere. So, it's a win for everyone. The thing that's also interesting, in particular this quarter, is that we also landed some large parts of hyperscalers. And hyperscalers typically have a culture of building everything themselves. And they certainly have the balance sheet and the human capital to support some of that build-out. Like if there was ever a set of companies for whom it makes sense to do it themselves, and we do those companies. And yet, we see that they have the same issue. When it comes to going as fast as they can, being as efficient as they can with their resources, like they come to us to replace some of the things that we're using before.

David M. Obstler

Chief Financial Officer

Two things, 2 metrics to look at that to make the points Oli, you're making, if you look at our platform adoption, and you see both the growth of the different categories and the extension of the categories out to lots of products that shows you that the consolidation on the Datadog platform has continued, and there's a very strong trend. And part of that is the movement solutions, as Oli mentioned, that are both open source, but also the competitive point solutions onto the platform. That's been a significant driver of the revenue growth for some time now, and that continued certainly in Q1.

Raimo Lenschow

Barclays Bank PLC, Research Division

Okay. Perfect. And then, David, for you, last year, and we did a lot of investments around go-to-market, especially on sales capacity. If you think about now the non-AI category doing better, how much of that is like people like the cloud migrations again. So that's like an industry trend and how much of that is like you guys actually being broader positioned?

David M. Obstler

Chief Financial Officer

Yes. It's a number of things, including one is the expansion of the platform, the consolidation, the successful ramping of sales capacity, which is not jeopardizing productivity and has resulted in ARR increasing and a good environment as well. And I think that's what we said last time, there are a number of factors. And certainly, what we're proving out here is the investments we've made in go-to-market and are continuing are paying off and we're the right decision. Oli, anything to add?

Olivier Pomel

Co-Founder, CEO & Director

Yes. And look, we, at the end of the day, there's clearly some market tailwinds with the adoption of AI and -- but also, we are outperforming all of our competitors at scale, and we're taking share, and that relates to the structural platform to where we expand with new products, the way these products are maturing and starting to win in their respective categories in the way we've successfully grown the sales capacity.

David M. Obstler

Chief Financial Officer

Certainly, the AI involvement trend has helped we're trying to do a separate that. AI investment is probably helping overall as well. But when you really take that out, you still see a very pronounced acceleration here. And that has to do with the factors that I mentioned and Oli talked about.

Operator

Our next question is coming from the line of Gabriela Borges of Goldman Sachs.

Gabriela Borges

Goldman Sachs Group, Inc., Research Division

Olivier, I find your comments on train versus inference, so interesting. Maybe just crystallize for us. Why do you think the training opportunity it's happening now or inflecting now? And then I had a for yourself for David, -- how do we think about the attach rate on trading versus inference of observability? Is there a way to benchmark observability spend as a percentage of inference spend, does that number change given the new data that you're seeing on the training site as well.

Olivier Pomel

Co-Founder, CEO & Director

So on the training side, training was very new a couple of years ago. It was something that was only done by very few companies, and it was in a way, very artisanal, like it was not a production workload. It was something that researchers were building, and that was very one-off and ongoing in ways. And now it's turning into production. It's turning into something that many more companies are doing. It's scaling by orders of magnitude, and it's becoming something that has to be on all the time, reliable and every minute you lose is or whether every fellow you have in your training around is a week you give away to the competition. And so, as a result, it becomes way more interesting as the market for a company like us. And we see some signs of that. Again, we didn't have a lot of it. We didn't see a lot of it last year. Now, all of a sudden, we're starting to see quite a bit of activity there and demand, then we have success landing with large customers with those products.

David M. Obstler

Chief Financial Officer

Yes. I think going back to the metrics that Oli talked about in terms of attach, we said that 6,500 customers are using our integrations and 20% of the customers and 80% of the ARR. So there is attach. I think it's earlier days for the training. That looks like it will be a contributor. But I think that's early, and I would sort of look at the larger attachment at this point as the evidence of inference, but also some training.

Operator

Our next question is coming from the line of Karl Keirstead UBS.

Karl Emil Keirstead

UBS Investment Bank, Research Division

Okay. Great. I wanted to start with Olivier and David, and congratulating all of you and the team on reaching that \$1 billion milestone well done. David, maybe the question is for you and to hone in specifically on the 2Q guide. Even if you put up a modest beat on that guide, it's going to be by order of magnitude, the largest sequential dollar at I think, in the company's history. And I just wanted to unpack what's giving you that confidence? And in particular, is there anything interesting to call out, David, in terms of the ramp of a couple of the larger research labs, one of which renewed with you guys in the fourth quarter, another one just landed. I presume they're ramping nicely in 2Q, but would love any color.

David M. Obstler

Chief Financial Officer

Yes. Let me unpack this in a couple of ways. As you know, we're recurring revenue model. So the biggest indication of the near-term of the next quarter is the ARR growth in the previous quarter. And when we said we had a record. So essentially, at the bedrock of this is sort of the run forward of ARR that we've already signed. The ARR add was very broad-based and was not very concentrated. So whereas we pointed out some very significant adds I would say that the first quarter and that ARR add was really diversified and from lots of different places. So, I think Oli will come in here, but the confidence that we have is you're right, we essentially take what we already have. We discount the growth trends that we've seen. And that produces what you exactly said, which is whatever your assumptions are on beat a very impressive sequential really due to what happened in Q1 and the rate of business accumulation by Datadog. Oli, do you want to add?

Olivier Pomel

Co-Founder, CEO & Director

Yes. I mean if you want to dive into what David just said, adds we broad-based. I mean look, when you look at why we had a great Q1, we also get customers in Q4. We talked about it a quarter ago. But even if you take out the customer we landed in Q4 that added the most revenue in Q1, we still had a record quarter in terms of ARR added. So, this is really broad-based. And we landed a few more customers in Q1 that don't contribute any revenue yet, but we expect to be big contributors in the future. So, when you put all that together, we feel very confident about Q2, hence the numbers you've seen.

Operator

And our next question will be coming from the line of Fatima Boolani of Citi.

Fatima Aslam Boolani

Citigroup Inc., Research Division

Oli, I wanted to double back on a question that was asked earlier with respect to telemetry volumes essentially going parabolic, and you are accessing a brand-new demand in the foray into training and monitoring and observing training model environment inside some of the world's largest frontier labs. And so I wanted to ask you about the structural changes to the capital intensity of the business. I mean your CapEx levels are still pretty respectable and pretty muted. So I wanted to get a better understanding of what sort of extrinsic or intrinsic engineering efforts you're undertaking to keep a very efficient CapEx envelope in spite of the fact that it seems like that would increase because of the torrent of telemetry we're seeing on the platform. And then as a related matter, we've seen a rise of sovereign data and data residency requirements kind of ramp as AI models move into the territory of national security and things like that. So just wondering if you can kind of talk to some of the engineering horsepower internally that you're leveraging to be able to keep a really tight command on capital intensity, and frankly, your gross margins?

Olivier Pomel

Co-Founder, CEO & Director

Yes, I mean, look, sort of the investments we're making right now, you we run most of our workloads on cloud, meaning you'll see all of that in OpEx, nothing CapEx. So. we have low CapEx. If it changes, we'll tell you, like if for some reason, we decide to make different kinds of investments and some of it more

front some it more CapEx, we'll tell you, but that's not the case today. We are definitely ramping up our investments in particular in R&D and in the scale of the models, we train ourselves and things like that. But right now, there's nothing that you can actually see in the numbers that move any needle but if that changes, also we'll tell you. We don't expect any change to So that's on the CapEx side. We are very different businesses in that way from the AI lab. On the subject of data residency and sovereignty of AI and things like that...we definitely see more push for more demand for that in the customer base. And for us, that means investment into areas One is in deploying into more geographies and having more certifications to sell to the public sector and to the highest level of the public sector. So, we mentioned today data center in the U.K., for example, and our certification, we're not stopping there in terms of the certification we're going after with a sell government. So that's an area of investment. Another area of investment is our bring you on cloud products and where we can actually run on our customers' infrastructure. And so we announced that, we read some products there, and we have heavy investment in that area. So we can support customers that want to operate in a slightly separate way from the rest of our customer base.

Operator

And our next question is coming from the line of Curt of Evercore.

Unknown Analyst

Congrats nice start. Oli, I was wondering if you could just give some thoughts on the idea of sort of security for agents. I think one of the big issues in terms of getting agents into production is sort of the security aspect of that? And how do you see Datadog plugging into that opportunity? And then just a quick one for David. Congrats on the FedRAMP reaching a milestone. Are your partner relationships in place to take advantage of this? I realize it will be a long-term opportunity, but just kind of curious how well established you are down there to start seeing some maybe bookings in that area.

Olivier Pomel

Co-Founder, CEO & Director

Yes. So on the security of agents, we interfere with that in 2 ways. So first, there's the agents will build ourselves because we are building a lot of automation inside of our products for our customers and agents that automatically identify but also resolve issues without you having to do anything. And there -- a lot of it has to do with understanding what permissions to apply, what kind of guardrails to apply, what kind of put to interface with the humans and how to make that the trust worthy and visible in the right way. And so that's pretty much the whole product surfaces to during data. The automation itself actually can work already. So you should expect to hear more about that at our conference. This is definitely one big area of investment for us. On the security aspects of our agents. Look, we believe in securities that you need to integrate, you can't just have point solutions that look at one sliver of the whole security posture. You need to look at everything all together. And that's one of the areas that we are also covering with our security efforts. So that's part of the whole platform action.

David M. Obstler

Chief Financial Officer

On the FedRAMP, we've been working on both the different certifications, but at the same time, we've been investing in the go-to-market function, both in terms of reps and channel partners for a number of years. Certainly, there's more investment to be done, but we invested ahead of the certifications because in this sector, building pipeline, et cetera, takes time. And certainly, channel partner relationships are a very important part of this. and we have been investing but also have more investment to do.

Operator

Our next question is coming from the line of Patrick Colville of Scotiabank.

Patrick Edwin Ronald Colville

Scotiabank Global Banking and Markets, Research Division

I guess, Olivier and David, you guys are very deliberate in your messaging on the prepared remarks. And I guess, I want to doublecheck the kind of wording of one of the comments. I think, David, you had a higher degree of conservatism to the largest customer. I guess, did I hear that right? And then does the higher degree of conservatism reference versus the other customer cohorts? Or does it reference versus your guidance philosophy in prior quarters vis-a-vis this customer?

David M. Obstler

Chief Financial Officer

It's both. It's the same guidance we used, and we're being very explicit. For all the business, except for the largest customers, we've always taken the drivers and discounting them. We -- for this particular customer, we took a higher degree of conservatism than the other part of the customer base. and discounted it more. And we were, I think, in the remarks and you interpret it correct, very explicit, and you're correct.

Olivier Pomel

Co-Founder, CEO & Director

I wouldn't give that much weight to do a very specific word. We deliberate but not all that deliberate. Similarly, both David and I have rusty voice today, but there's a man.

David M. Obstler

Chief Financial Officer

But I will remind everybody we did not change. So if the question also is did we change? Or is this a different methodology of both the overall and the large customer then the guidance the last quarter or the previous? The answer is no. It's the same methodology that we've had. So, no change, but that has been what we've always been doing.

Patrick Edwin Ronald Colville

Scotiabank Global Banking and Markets, Research Division

Okay. And Olivier, can I ask about your comments about the hyperscalers because I thought that was particularly interesting. And the reason why is, I don't think you called them out previously before, and they are so prevalent in the modern tech stack. To your point, they could do this themselves. So, I guess how are they using Datadog? Is it for more kind of traditional observability? Or is it for these newer areas like GPU monitoring that Datadog has performed so well of late.

Olivier Pomel

Co-Founder, CEO & Director

Well, it's both actually. When you look in general at large AI customers, they use Datadog at the way other companies are largely with a fairly broad set of our products to cover the full circuit of liability. What's new is we now have a product for GPU monitor. It's a very new product. And we see the hyperscalers that are coming to us for training workloads in particular being very interested in that. So again, it's too early in the product life cycle and the customer life cycle for these specific customers to go definitive victory there, but we see that as a very encouraging sign of where the market might go in the future. Because we think this might be a bellwether of what the next 10, 100, 500 companies that are going to start training workloads are going to want to do. We have some signs that go beyond the customers we signed this quarter that point that way too.

Operator

And our next question is coming from the line of Peter Weed of Bernstein Research.

Peter Weed

Bernstein Institutional Services LLC, Research Division

And I'll echo others on the momentum. Great to see one of, I think, the great successes you talked about was landing a couple of the AI labs for the hyperscalers. Although I think on the other hand, you've talked in the past around hyperscalers are typically building observability in-house. What is it really about the AI workloads that are making it more attractive for them to use Datadog? And what might give you confidence that Datadog might be more persistent with them in these types of workloads and that's kind of a signal for maybe how other customers might use Datadog around AI differentiated from things that they might be able to bring in house to other places?

Olivier Pomel

Co-Founder, CEO & Director

The same as all our customers. It's high stakes, high complexity and not core. They have to be most differentiated. They can't afford to be late, and it's a really hard job to do that. So, it's what we built our whole business on, and it's also very true at the highest level for the largest companies.

Peter Weed

Bernstein Institutional Services LLC, Research Division

Yes. No, I was just going to say it. But I guess, the point is you've emphasized that those largest customers have been able to go in-house on some other things. Is there something unique about AI that prevents them from doing that here?

Olivier Pomel

Co-Founder, CEO & Director

Well, I think the urgency of their development efforts focuses the mines. That's what I would put it. I would say, it forces you to figure out what's core and what's not core and what you need to do to maximize your chances of success. And again, it's the same thinking all of our customers have all the time. I think the equation for hyperscalers has often been, say, different because they have, let's call it, unlimited access to staffing. And they sort of set their own time horizons for the developments they wanted to make. I think the situation is a little bit different with the AI labs maybe.

Operator

The line of Gregg Moskowitz of Mizuho.

Gregg Steven Moskowitz

Mizuho Securities USA LLC, Research Division

And I'll add my congratulations on a terrific quarter. Just one for me. Oli, I know it's not GA yet, but curious if you have any early feedback on your new cloud prem offering. As you noted earlier, providing the ability for Datadog to run on customer infrastructure. Could this be another yet another, I should say, incremental growth opportunity for Datadog? What are your expectations for this?

Olivier Pomel

Co-Founder, CEO & Director

Well, definitely, we think -- I think there was a question earlier on data residency and leaving customers environment, we definitely see a great opportunity there. It is chance that a good portion of the market means this way in the future. Today, it's not the largest part of the market, but we definitely see a potential for that. So, we're investing heavily in that sort of our product. We're trying to see some interesting customer traction there. So we think this can be another growth lever differently. We also think that it can help us getting into some extremely large-scale workload where customers would not have considered SaaS offering before, where we can be in the running. So that's very exciting. All right. And I think that was our last question. So I want to thank you all for attending the call. And I remind you that we have our conference in just a bit more than a month, and I hope to see many of you there. So thank you all.

Operator

This concludes today's program. You may all disconnect.