

**Datadog, Inc. NasdaqGS:DDOG**  
**FQ1 2025 Earnings Call Transcripts**  
**Tuesday, May 6, 2025 12:00 PM GMT**

# Table of Contents

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

# Call Participants

## EXECUTIVES

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

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

**Yuka Broderick**  
*Investor Relations*

## ANALYSTS

**Andrew Michael Sherman**  
*TD Cowen, Research Division*

**Brad Robert Reback**  
*Stifel, Nicolaus & Company,  
Incorporated, Research Division*

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

**Gregg Steven Moskowitz**  
*Mizuho Securities USA LLC, Research  
Division*

**Howard Ma**  
*Guggenheim Securities, LLC, Research  
Division*

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

**Kasthuri Gopalan Rangan**  
*Goldman Sachs Group, Inc., Research  
Division*

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

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

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

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

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

# Presentation

## Operator

Good day, and thank you for standing by. Welcome to the Q1 2025 Datadog Earnings Conference Call. [Operator Instructions] Please be advised that today's conference is being recorded.

I would now like to hand the conference over to your first speaker today, Yuka Broderick, SVP of Investor Relations. Please go ahead.

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

Thank you, Antoine. Good morning, and thank you for joining us to review Datadog's first quarter 2025 financial results, which we announced in our press release issued this morning. Joining me on the call today are Olivier Pomel, Datadog's Co-Founder and CEO; and David Obstler, Datadog's CFO.

During this call, we will make forward-looking statements, including statements related to our future financial performance, our outlook for the second year and the fiscal year 2025 and related notes and assumptions, our gross margins and operating margins, our product capabilities, our ability to capitalize on market opportunities and usage trends.

The words anticipate, believe, continue, estimate, expect, intend, will and similar expressions are intended to identify forward-looking statements or similar indications of future expectations. These statements reflect our views only as of today and are subject to a variety of risks and uncertainties that could cause actual results to differ materially.

For a discussion of the material risks and other important factors that could affect our actual results, please refer to the Form 10-K for the quarter ended December 31, 2024. Additional information will be made available in our upcoming Form 10-Q for the fiscal quarter ended March 31, 2025, and other filings with the SEC. This information is also available on the Investor Relations section of our website, along with a replay of this call. We will also 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](mailto:investors@datadoghq.com).

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

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

Thanks, Yuka, and thank you all for joining us this morning to go through our Q1 results and what was a solid start to the year. Let me begin with a review of our Q1 financial performance.

Revenue was \$762 million, an increase of 25% year-over-year and above the high end of our guidance range. We ended with about 30,500 customers, up from about 28,000 a year ago. We ended Q1 with about 3,770 customers with an ARR of \$100,000 or more, up from about 3,340 a year ago. These customers generated about 88% of our ARR. And we generated free cash flow of \$244 million, with a free cash flow margin of 32%.

Turning to platform adoption. Our platform strategy continues to resonate in the market. At the end of Q1, 83% of customers were using 2 or more products, up from 82% a year ago. 51% of customers were using 4 or more products, up from 47% a year ago. 28% of our customers were using 6 or more products, up from 23% a year ago. And 13% of our customers were using 8 or more products, up from 10% a year ago.

We are pleased to see that customers are adopting more products, and I'd like to highlight 2 of our newer products with you. First, Flex Logs is off to a fast start and now exceeds \$50 million in ARR. Flex Logs has achieved this milestone in 6 quarters, which is the fastest ramp we've seen to that level, and which echoes its value to customers as well as the size of the log management market opportunity.

I'll also note that by adopting Flex Logs, our customers are adding new use cases at the right economics. And these Flex Logs adopters ultimately spend more on Datadog log management as well as more on our overall platform.

Second, our Database Monitoring product is approaching the \$50 million ARR level as well and is growing 60% year-over-year. Database Monitoring has now been adopted by over 5,000 customers. We are very excited about the early traction we're seeing there,

and are doubling down on our investment into broader Datadog capability as we see strong demand signals in that area, and we'll come back to that in a few minutes.

Now let's discuss this quarter's business drivers. Overall, we saw trends for usage growth from existing customers in Q1 that were in line with our expectations. We are seeing high growth in our AI cohort as well as consistent and stable growth in the rest of the business. We also had a strong bookings quarter, with particularly strong execution in new logos and larger bookings.

Dollar bookings for new logos were up over 70% year-over-year and much stronger than our typical seasonal softness in Q1. And on the large deal side, in Q1, we signed a total of 11 deals with a TCV of \$10 million or more, up from just \$1 million in the year ago quarter, as we continue to expand our business with larger customers. 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.

Now moving on to R&D. We continue to see rising customer for next-gen AI observability and analysis. At the end of Q1, more than 4,000 customers used one or more Datadog AI integrations, and this number has doubled year-over-year. With LLM Observability, we are seeing continued growth in customers and usage as they seek to manage end-to-end model performance, security and quality.

I'll call out the fact that the number of companies using LLM Observability has more than doubled in the past 6 months. And we are adding to Bits AI, with capabilities for customers to take action with workflow automation and App Builder, using next GenAI to help our customers remediate issues more quickly and move towards auto remediation in the future. Zooming out, we're making progress on all of our AI initiatives, and you should expect more announcements in this area at DASH, our user conference taking place in June.

Moving to security. Our teams have been very busy building our products and features for our customers' DevSecOps. To give you a quick overview of our capabilities. First, we have a comprehensive set of products to identify and manage vulnerabilities across software and infrastructure.

In infrastructure, our cloud security product identifies vulnerabilities in hosts, containers, Kubernetes clusters and infrastructure as a code. Our security customers can use agentless scanning to cover their entire infrastructure stack in minutes. And existing Datadog customers using our lightweight agent immediately gain deep, granular and timely security visibility.

On the application vulnerability side, our code security product identifies vulnerabilities in code from development to production, and for first-party code as well as third-party open-source libraries. This product area has launched very recently and already has over 1,000 customers paying for the product. Because we bring visibility to production workloads, we are uniquely positioned to identify which vulnerabilities are most critical in production and break down silos between developers, DevOps and security teams.

Second, in security. As vulnerabilities face threats and attacks, our threat management product helps our customers identify and remediate them. They can use our Cloud SIEM to identify threat logs, and they can further protect from threats in infrastructure with workflow protection; and in software, with app and API protection.

Finally, our customers use our Sensitive Data Scanner product to discover, classify and redact sensitive data at scale across their logs, traces, events, user sessions, data source code and all the way to LLM prompts. But we have much more to do. Today, we are serving over 7,500 customers with our security products or about 1/4 of our total customer base. And over half of our Fortune 500 customers use our security products, a good sign of our opportunities with the largest enterprises.

Now moving on from security. Last month, we announced our plans to launch our latest data center in Australia. We see a large opportunity to serve our Australian customers and help them meet local data residency, privacy and security requirements.

Finally, we really announced a couple of acquisitions. First, we acquired Eppo, a next-generation feature management and experimentation platform. The Eppo platform helps increase the velocity of releases, while also lowering risk by helping customers to release and validate features in a controlled manner.

Eppo augments our efforts in product analytics, helping customers improve the variance and tie feature performance to business outcomes. More broadly, we see automated experimentation as a key part of modern application development, with the rapid adoption of AI generative code, as well as more and more of the application logic itself being implemented with nondeterministic AI models.

Second, we also acquired Metaplane, a data observability platform built for modern data teams. Metaplane helps prevent, detect and resolve data availability and quality issues across the company's data warehouses and data pipelines. We've seen for several years now that data freshness and quality were critical for applications and business analytics. And we believe that they are becoming key enablers of the creation of new enterprise AI workloads, which is why we intend to integrate the Metaplane capabilities into our end-to-end data observability offerings.

We are very excited to welcome both the Metaplane and the Eppo teams to Datadog as we have a lot to build together, and that fits for our product engineering. Our teams are very hard at work this quarter, and we're looking forward to sharing many new products and feature announcements at our DASH user conference on June 10 and 11 in New York City.

Now let's move on to sales and marketing. As I mentioned earlier, we have a number of great new logo wins and customer expansions this quarter, so let's go through a few of those.

First, we landed a 7-figure annualized deal with one of the largest U.S. car manufacturers. This customer has a complex hybrid environment, including on-prem, multiple clouds, in-car IoT and mobile apps. They expect to unify observability across teams and across all their tech stacks, while accelerating root cause analysis. And they are starting with 13 Datadog products, consolidating a dozen tools and rolling out to dozens of business units.

Next, we landed a 7-figure annualized deal with a major Latin American bank. They expect to use our unified observability to reduce operational costs and enable autonomy for departments that previously had to depend on specialized teams for visibility. This customer is starting with 6 Datadog products and is replacing 3 existing tools.

Next, we landed a 7-figure annualized deal with a major American pet supplies company. These customers struggle with tools pro and limited user option. With Datadog, they expect to save over \$1 million every year, both in engineering time and avoidance of lost revenue. This customer is starting with 11 Datadog products, including Cloud SIEM, and is replacing 7 commercial tools.

Next, we welcome back an insurance tech customer with a 6-figure annualized deal. This customer found that their previous observability tool involved manual workflows and customization, a high operational overhead and user frustration and fatigue. By returning to Datadog, they expect to benefit from Datadog's ease of use and out-of-the-box capabilities, while using our built-in usage controls to manage observability and cloud costs. This customer now expects to use Flex Logs, Cloud Cost Management and OnCall, among the 10 products they plan to adopt.

Next, we signed a 7-figure annualized expansion with one of the largest U.S. health insurers. This customer is using Datadog across dozens of business units to support millions of customers. More recently, they have been using Datadog to significantly improve customer experience and reduce time-consuming and expensive outages. As an example, one team estimated reductions in mean time to resolution from 3 to 4 hours, down to 3 to 4 minutes by using Datadog. With this expansion, this customer is using 17 products in the Datadog platform, including the full Datadog security suite.

Finally, we signed a 7-figure expansion as an annualized contract with a leading next GenAI company. This customer needs to reduce tool fragmentation to keep on top of its hyper growth in usage and employee headcount. With this expansion, the customer will use 5 Datadog products and will replace commercial tool for APM and log management. And that's it for another productive quarter from our go-to-market teams.

Before I turn it over to David for a financial review, let's have a few words on our longer-term outlook. We recognize that there are many cross currents impacting the global economy right now. But our view of our long-term market opportunities remains unchanged. We continue to believe digital transformation and cloud migration are long-term secular growth drivers of our business, as well as critical for every company to deliver value and gain competitive advantage. And we continue to focus on delivering innovation and value to our customers against their mission-critical needs, including their AI efforts.

Now more than ever, we feel ideally positioned to help customers of every size and in every industry to transform, innovate and drive value for technology adoption.

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

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

Thanks, Olivier. Q1 revenue was \$762 million, up 25% year-over-year and up 3% quarter-over-quarter. To dive into some of the drivers of Q1 revenue growth. First, overall, we saw trends for usage growth from existing customers in line with our expectations and similar to the second half of 2024. We saw a continued rise in contribution from AI-native customers who represented about 8.5% of Q1 ARR, up from about 6% of ARR last quarter and up from about 3.5% of ARR in the year ago quarter.

AI-native customers contributed about 6 points of year-over-year revenue growth in Q1 versus about 5 points last quarter and about 2 points in the year ago quarter. We continue to believe that adoption of AI will benefit Datadog in the long term, but we remain mindful that we may see volatility in our revenue growth on the backdrop of long-term volume growth from this cohort as customers renew with us on different terms and as they may choose to optimize cloud and observability usage.

Next, regarding usage growth by customer segment. Year-over-year usage growth with our enterprise customers remain healthy. It's a bit lower than last quarter, which we see as a product of the volatility that can occur among customers from quarter-to-quarter. Meanwhile, we saw strong booking activity from our enterprise customers in Q1. And as Olivier noted, this included some large TCV deals.

From our SMB and mid-market customers, excluding the AI cohort, year-over-year usage growth was roughly similar compared to last quarter. As a reminder, we define enterprise as customers with 5,000 employees or more, mid-market as customers with 1,000 to 5,000 employees, and SMB as customers with less than 1,000 employees.

Looking ahead to April, our usage growth was consistent with the year-to-date trends. As usual, we have contemplated near-term trends in our guidance. Regarding retention metrics, our trailing 12-month net revenue retention percentage was in the high 110% in Q1, similar to last quarter. Finally, our trailing 12-month gross retention revenue percentage remained stable in the mid- to high 90s.

On new logos, gross new logo additions were roughly the same as in Q1 last year, but dollar new logos increased 70% year-over-year, indicating a higher average land in the SMB, mid-market and enterprise sectors. And our pipeline for Q2 is strong and growing healthily year-over-year. As a reminder, our sales pipeline doesn't convert into revenue immediately.

Now moving on to our financial results. Billings were \$748 million, up 21% year-over-year. Remaining performance obligations or RPO, was \$2.31 billion, up 33% year-over-year and current RPO growth was about 30% year-over-year. RPO duration was roughly flat year-over-year. We continue to believe revenue is a better indication of our business trends than billings or RPO as those can fluctuate relative to revenue based on the timing of invoicing and the duration of customer contracts.

Now let's review some of the key income statement results. Unless otherwise noted, all metrics are non-GAAP. We have provided a reconciliation of GAAP to non-GAAP financials in our earnings release.

First, gross profit in the quarter was \$612 million and gross margins was 80.3%. This compares to a gross margin of 81.7% in the last quarter and 83.3% in the year ago quarter. While gross margin remained in the range that we have expected over the long term, our cloud hosting costs rose more quickly than we expected in Q1, as we supported large growth spikes from some of our largest customers.

We also continue to innovate with new products and capabilities for our customers, which tend to put downward pressure on gross margins in the short term. While we expect some of the costs to support customers to persist, we are also focused on executing projects to improve our cloud cost efficiency and expect to realize savings throughout the rest of the year.

Our Q1 OpEx grew 29% year-over-year, similar to the 30% last quarter and roughly as expected as we continue to execute on our hiring plans. As we have spoken about in previous quarters, it has been our plan to grow our investments to pursue our long-term growth opportunities. And we've been successful in increasing sales rep headcount with over 25% year-over-year growth in total reps, including over 30% growth year-over-year in enterprise reps.

This investment has been weighted a little bit towards international expansion, where sales rep headcount growth was in the mid-30s percent year-over-year. We believe that there is white space to support this growth and that we will keep seeing results as this capacity ramps as we have in the past.

In addition, we grew our R&D headcount by over 30% year-over-year, with R&D expense as a percent of sales at 30% in Q1, so can we deliver on the rapid pace of innovation for our customers. As always, we continue to prioritize our investments to balance near-term adjustments with our long-term plans and continue to look for opportunities to optimize our operating costs. Q1 operating income was \$167 million for a 22% operating margin compared to 24% last quarter and 27% in the year ago quarter.

Now turning to the balance sheet and cash flow statements. We ended the quarter with \$4.4 billion in cash, cash equivalents and marketable securities. Cash flow from operations was \$272 million in the quarter. And after taking into consideration CapEx capitalized software, free cash flow was \$244 million for a free cash flow margin of 32%.

And now for our outlook for the second quarter and the full fiscal year 2025. First, our guidance philosophy overall remains unchanged. As a reminder, we based our guidance on trends observed in recent months and apply conservatism on these growth trends.

So for the second quarter, we expect revenue to be in the range of \$787 million to \$791 million, which represents 22% to 23% year-over-year growth. Non-GAAP operating income is expected to be in the range of \$148 million to \$152 million, which implies an operating margin of 19%. As a reminder, in Q2, we will be hosting our DASH user conference, which we estimate to cost about \$13 million and which we have reflected in our operating income guidance.

Non-GAAP net income per share is expected to be \$0.40 to \$0.42 per share based on approximately 361 million weighted average diluted shares outstanding. To note, weighted average diluted share count is expected to decline sequentially as the share related to the 2025 convertible note will be removed upon redemption.

For fiscal year 2025, in total, we expect revenue to be in the range of \$3.215 billion to \$3.235 billion, which represents a growth rate of 20% to 21%. Now we have raised our 2025 revenue guidance range by \$40 million related to the previous guidance, which incorporates higher revenues in the first half of 2025 based on our Q1 results and our visibility as of today into Q2. Our implied guidance in the second half of 2025 is roughly unchanged.

Non-GAAP operating income is expected to be in the range of \$625 million to \$645 million, which implies an operating margin of 19% to 20%, relative to our previous operating income guidance range of \$655 million to \$675 million. The difference is mainly due to lower gross profit as a result of the previously discussed lower gross profit margin, offset by higher revenues.

As we said before, we're focused on executing cost efficiencies in our cloud costs and believe our gross margins will remain in our historical range. Overall, plans for OpEx investment are roughly unchanged, with continued investment in hiring across R&D and sales and marketing.

There are some changes with distribution within that. We expect \$15 million of higher international costs due to currency rate changes and \$10 million in net expected additional costs from our recently announced acquisition, offset by other OpEx savings. Non-GAAP net income per share is expected to be in the range of \$1.67 to \$1.71 per share based on approximately 362 million weighted average diluted shares outstanding.

And finally, some additional notes on the guidance. The remaining approximately \$635 million principal of our 2025 convertible notes will mature in June and we expect to redeem this mainly in cash. We estimate that the GAAP purchase price from our Q2 acquisition activity will be about \$180 million, of which we estimate about \$110 million to be paid in cash during Q2.

We expect net interest and other income for the fiscal year 2025 to be approximately \$140 million, and we expect cash taxes to be about 1% of revenue or about \$30 million to \$35 million. We continue to apply a 21% non-GAAP tax rate for 2025 and going forward. And finally, we continue to expect capital expenditures and capitalized software together to be in the 4% to 5% of revenue range in the year.

Finally, to summarize, we are pleased with our execution in Q1. We are well positioned to help our existing and prospective clients with their cloud migration and digital transformation journeys. And I want to thank Datadogs worldwide for their efforts. Now 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 comes from Mark Murphy from JPMorgan.

### Mark Ronald Murphy

*JPMorgan Chase & Co, Research Division*

Congratulations on a great performance. Olivier, we noticed the CEO of Anthropic recently said that within 12 to 18 months, that AI is going to be writing 100% of all code. I'm sure there's a bit of hyperbole there, but directionally, it's intriguing. Can you comment on the opportunity that might open up for Datadog? If that sheer volume of applications being put into production starts to rise because AI writes so much code so rapidly. And just does that AI-generated code require more or less monitoring than human-written code? And then I have a quick follow-up for David.

### Olivier Pomel

*Co-Founder, CEO & Director*

Yes, that's a great question. And there is definitely a big transition that is happening right now, like we see the rise of AI written code. We see it across our customers. We also see it inside of Datadog, where we've had very rapid adoption of this technology as well. While I don't think this is going to replace all our software engineering, and I'm pretty sure that Anthropic is still hiring software engineers too, I do expect big changes to come to the way software is being shipped and being run this way.

The way we see it is that it means that there's a lot less value in writing the code itself. Like everybody can do it pretty quickly, can do a lot of it. You can have the machine to do a lot of it, and you complement it with a little bit of your own work. But the real difficulty is in validating that code, making sure that it's safe, making sure it runs well, that it's performing and that it does what it's supposed to do for the business. Also making sure that when 15 different people are changing the code at the same time, all of these different changes come together and work the right way, and you understand the way the different pieces interact in a way.

So the way we see it is this moves a lot of the value from writing the code to observing it and understanding it in production environments, which is what we do. So a lot of the investments we're making right now, including some of the acquisitions we've announced build towards that, and making sure that we're in the right spot. So we can tell you exactly what every piece of code you've written, make sure that it works well, you understand it well, and it does what it's supposed to do for the business.

### Mark Ronald Murphy

*JPMorgan Chase & Co, Research Division*

Okay. Okay. Very interesting. And then David, the -- it's just amazing, the booking stats, so impressive, 11 deals, over \$10 million, only one a year ago. And then these big AI numbers, the combination there, it's just incredible. What do you attribute that level of booking strength to? It's a quarter where you had the specter of a trade war and it seemed like it was weighing on business confidence. So I'm wondering if you can comment on the booking strength.

### David M. Obstler

*Chief Financial Officer*

Yes. We entered the quarter with a strong pipeline and I think it's consistent with our investments we're making in go-to-market, where we're increasing our investments in enterprise and across the board. And so -- and we've been good about putting value to clients across the platform, allowing us to land larger, as you saw, consolidate and add more value. I think that much of this is related to product strength and expansion of quota capacity.

### Olivier Pomel

*Co-Founder, CEO & Director*

Yes. And look, we -- like everybody, we read the news and we see -- we look hard for any signs of trouble in our customer base. We haven't seen that in the dealmaking so far. We -- our sales cycles haven't been affected. Our pipelines are growing healthily.

And I can tell you that at least one of the customers we mentioned in the call, one of the new lands, is also a company that is very affected by the tariffs and is having to moderate its own plans for the future. I think what this tells you is that cloud migration and

observability in particular are deflationary for these communities and they help -- these are tools that help them save money and move faster, and we're on the right side basically, of the problem for them.

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

And to echo a little bit what Oli said in our prepared remarks, we continue to see that. We said we have, in Q2, same quality of pipeline higher than last year. So it continues -- we continue to have that situation.

**Operator**

Our next question comes from Sanjit Singh from Morgan Stanley.

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

Congrats on the quarter. I want to pick up where Mark left off. And as we talk about kind of the core drivers, digital transformation, cloud migrations, what are the trend lines on the cloud migration side? It seems to -- based on your comment, Oli, it sounds like that's likely -- or could actually pick up potentially going into a lower economic environment. So in terms of just what we've seen through the early parts of the year, what are the trend lines on core cloud migration?

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

So it's very consistent. It's consistent with what we've seen before. It's also consistent with what you've heard from the hyperscalers over the past couple of weeks. So I would say it's steady, unremarkable, it's not really trending up, not trending down right now, but we see the same desire from customers to move more into the cloud and to lay the groundwork so they can also adopt AI, because digital transformation and cloud migrations are prerequisites for that.

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

Awesome. I want to talk a little bit about some of the expansion -- additional expansion opportunities into data observability. I mean that's been a space that's been talked about, sort of buzzy the last couple of years. I'm not sure it's gotten super big just yet. Maybe just some of the -- what the vision is about moving into data observability and how consequential of an opportunity it could be for Datadog.

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

So it's a space we've been watching for a while. Initially, we were worried that this was maybe a little bit too small in impact because if we're just talking about debugging the availability of reports for executives, like it is definitely valuable, but we thought it might be like a smaller market opportunity in the end.

What we see, though, is that the field is evolving into a big enabler or it can be the opposite of an enabler, if you don't do it right, for building enterprise workloads -- for building AI workloads, sorry. So in other words, making sure the data is being extracted from the right place, transformed the right way and is being fed into the right AI models on the other end. And so we're building towards that with this acquisition of Metaplane we just announced.

And just to help you feed that into what we're doing, we already had some building blocks for data observability. We built a Data Streams Monitoring product for streaming data that comes out of Q such as Kafka, for example. We built a data jobs monitoring product that monitors spark jobs and large data transformation jobs. We have a database monitoring product that looks at the way you optimize queries and optimize database performance and cost.

And by adding data quality and data pipelines with Metaplane, we have a full suite basically that allows our customers to manage everything from getting the data from their core data stores into all of the products and AI workloads and reports they need to populate that data. And so we think it's a big opportunity for us.

**Operator**

Our next question comes from Raimo Lenschow from Barclays.

**Raimo Lenschow**

*Barclays Bank PLC, Research Division*

Perfect. Congrats from me as well. Two quick questions, actually more for David this time. David, if you think about the guidance, I know you've been -- you've always been a conservative team. You raised it more than the Q1 beat. So obviously, you're kind of seeing good trends in Q2 there as well. What was the thinking process in terms of kind of showing us that upside that you kind of might expect in Q2 versus kind of buffering that and just kind of given the uncertain environment? It's the first question.

Second question is on the gross margin change for the full year guidance. Can you just remind us like, obviously, that's 3 months after you guided the first time. Can you remind us a little bit like what you saw there and what's the action that you're taking there to kind of change that?

**David M. Obstler**

*Chief Financial Officer*

Yes. So on the first question, we haven't changed any of our strategy towards guidance. We look at the recent trends. I think it's important to note that we obviously beat in Q1, and we sort of looked at the run rate and still discounted the current use growth rate. So still put the conservatism into Q2, and this was the result of the guide. So it indicates the pace of business and where we are, but still discounts the trends as we always do.

And another thing that's important to note is because of the uncertainty in the market, we left the second half of the year unchanged. We essentially have done that. There's no change in our strategies that we do all the time. But as we go out more in time, we have less visibility. So I think there's no change in philosophy, it's just sort of a follow-through on our run rate.

In terms of gross margin, as a reminder, we've always said that our gross margins would fluctuate within the range. And we've been in the range of the upper 70s to low 80s for a long time. And so what we -- we sort of designed the business that way. At sometimes, we're going to lean into investment. And in some cases, we'll optimize.

In Q1, we -- as we talked about, we did lean into investment, and we also had a little more of a spiky pattern from our customers, which we're learning about. We're learning as we grow in size, how to deal with that. And so from those learnings, we feel like on that, we can do a better job in provisioning. And as we always said, we run in the spectrum.

So I think there were higher cloud costs than we expected. It is largely due to those 2 factors. And we said some of that will persist and some of it, we will work on optimizing. This is no different than on that sort of pattern that we have over time of growth and then optimization. Oli?

**Olivier Pomel**

*Co-Founder, CEO & Director*

Yes. And just to echo that, we have the same engineers basically shipping functionality and optimizing performance. And we -- there's a certain range of margin we're comfortable with. I would say last year, we were very comfortable with where we were. Now we're getting a little bit less comfortable with what we've seen over the last quarter. And so we're shifting some resources from -- towards optimization so we can right that.

As David said, in addition to shipping a lot of new functionality that might not be as optimized out of the gate as the rest in terms of the system utilization, we also saw some spiky growth from some of our largest customers. And that might continue and especially as they deploy new workload. But our job is to optimize, and we feel very confident that we can remain in the range we've been in before for gross margin or even better in the future. There's no question about that.

**Operator**

Our next question comes from Kash Rangan from Goldman Sachs.

**Kasthuri Gopalan Rangan**

*Goldman Sachs Group, Inc., Research Division*

Oli, one for you. And one for you, David. Oli, when you look at the AI market, certainly it looks like we're about to move from training into inference. Where does that leave Datadog from a product perspective where you can add more consistent value without being optimized? And one for you, David. As you look at the stepped-up research development, sales and marketing investments, how from a financial perspective are we to see the incremental benefit of those incremental investments?

**Olivier Pomel**

*Co-Founder, CEO & Director*

Yes. So on the workloads turning more towards inference, so there's definitely more product to build there. So we have a -- so we built an LLM Observability product that is being -- that is getting increasing usage from customers as they move into production. And we think there's more that we need to build both down the stack closer to the GPUs and up the stack closer to the agents that are being built on top of these models. So you should expect to hear more from us at our conference on those topics because there's, I think, quite a bit to -- that is emerging as a set of needs from customers as they move towards inference.

**David M. Obstler**

*Chief Financial Officer*

And on the investments, I think you're seeing it, we're trying to report on it regularly when we're making disclosures of being over \$50 million -- or being around \$50 million in database and then in Flex. That is the evidence of being able to monetize the investments in R&D and continue to do that in the various product offerings. So that's how we look at that, and we've gotten a lot of evidence over the years and continue to report on that.

We have always said that the investment for R&D takes 2, sometimes 3 years to realize itself because of building in products, but we're reporting on investments that we made previously and the result of that in products. In sales and marketing, it's the same thing where we identify end markets or customers, which we have covered, feel like we can add additional capacity. And there, we have to go through the ramp period. So the increase of capacity that we've been executing on, we said we'll pay back in sort of the year plus as we ramp our reps and get returns, and we'll continue to report on that.

**Olivier Pomel**

*Co-Founder, CEO & Director*

Yes. And just as a quick anecdote, like we -- so we talked about DB Database Monitoring as a product that has taken off. And that's really a product that shows the power of the platform. So we have a product that is over \$50 million -- or closing on \$50 million in ARR with a large number of customers. And that's a product that was built basically by 10 to 15 people. And that's just the ability for a small team to build on top of our platform and on top of all of the existing services we have that already collect customer data and deliver the data to them and to make a large difference in terms of revenue in a short amount of time.

**Operator**

Our next question comes from Jake Roberge from William Blair.

**Jacob Roberge**

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

Could you talk a little bit more about what you're seeing with the AI-native cohort? It sounds like growth remains solid in that segment, but still some potential for optimization there. So could you talk about what you're seeing from some of those AI-native contracts that have already come up for renewal and just how those conversations have been trending?

**Olivier Pomel**

*Co-Founder, CEO & Director*

Yes. So I mean all the contracts that come up for renewal, they are healthy. The trick with the cohort is that it's growing fast. There's also a revenue concentration there. We now have our largest customer in the cohort, and they're growing very fast. And on the flip side of that, we also have a larger number of large customers that are also growing. So we -- I think we mentioned more than 10 customers now that are spending \$1 million or more with us in that AI-native cohort and that are also growing fast.

So when you add up all of that, like the number of newer customers that are growing fast and some concentration in one customer is larger than the others, like there is definitely some potential volatility there, and we want to be careful. We've seen that movie before. We've seen that with the cloud native a few years ago as they were growing very, very fast out of COVID, and they had to optimize after that quite sharply.

Now of course, the situation is different, like it's a small part of our business today, whereas a very large part of our business was cloud native a few years ago, but we want to be mindful of what might happen next. And that's also part of how we incorporate that revenue and its growth in terms of our guidance for the rest of the year.

**Jacob Roberge**

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

Okay. That's helpful. And then great to see the strong ramp in Flex Logs over the past year. Can you talk a little bit more about what you're seeing in the log management market specifically? We've definitely been hearing about some more disruption at the start of the year from the recent acquisitions in the space. So just curious if you're seeing an even more meaningful opportunity on that front this year.

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

We see tremendous opportunity there. We validated that Flex Logs was really solving a big problem for customers. We've also validated the fact that this was net new use cases for us. And this was not really just a cheaper product that cannibalizes the rest of what we have, like we do see customers adopting Flex Logs and then consuming a lot more of everything after that. So this is great for us.

We also see a ton of opportunity to displace some existing players in log management. We've built a lot of functionality towards that to close whatever gap in functionality or specific implementation of functionality that might have been perceived by customers on that end. And so we see a ton of opportunity there.

And one of the reasons we're growing the sales capacity in the way we are currently is that we see those opportunities in the market. We see great return on investment when we deploy more sales capacity to go after them because these tend to be large enterprise opportunities that require some sales capacity, and we intend to fully capitalize on that. So we're super bullish about this.

**Operator**

Our next question comes from Gregg Moskowitz from Mizuho.

**Gregg Steven Moskowitz**  
*Mizuho Securities USA LLC, Research Division*

I have 2. I'll just ask them concurrently. First, the AI-native cohorts, getting back to that. So the percentage of total ARR coming from this cohort rose really impressively. Just so that we have a little better understanding of this dynamic because Oli, you just mentioned or referenced customer concentration. Can you say if the increase from a percentage basis was driven by perhaps 1 or 2 of these larger customers or was it broader based than that?

And then secondly, for David, just given the second half step-up in sales and marketing investment that we saw last year at Datadog that seems to have continued into 2025. How is your productivity tracking at this stage compared to your expectations?

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

Yes. So I mean, look, on the AI side, we do have, as I mentioned, one customer large and the others there, they're contributing more of the new revenue than the others. But we see growth in the rest of the cohort as well. So again, it's fairly typical.

We've always had in every single one of our fast-growth cohort, we do have some form of revenue concentration. And I remind you that a small fraction of our customers are the ones that pay us more than \$100,000 of ARR, and they represent almost 90% of our revenue. So that's sort of how the business is structured. But we do -- we are careful about the size and the rate of growth of those customers there, which is also why we call attention to potential volatility there in the future.

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

Yes, that's something we watch a lot. So once the reps are ramped, we look at productivity and do we have productivity from those added that are similar? And the answer is yes. It is due to where we're putting the reps. We're not doing it sort of in a nontargeted way. We're looking at territories. So when we see ramps -- reps ramp, we see similar productivity. And that gives us the signal that we are doing the right thing, and we can continue to add additional reps.

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

Just as a quick reminder, the -- when you think of growing capacity, I think it's a little bit different for sales than it is from engineering. For engineering, you can do it in a top-down fashion. You can decide how much you're going to grow the team and then you can distribute your engineering to different types of products after that, more or less.

On the sales side, you need to do it in the bottom-up fashion because you need to understand which territories can carry how many new reps, which ones have the right productivities, which one have the right market opportunity and where you're going to put people. So it's not as easy as saying we're going to grow the team by 30%. You actually need to figure out who's going to go where and how you're going to grow this team and how you're going to make everyone successful.

So we've done the hard work for doing that. We've done it for many years now. And the growth in capacity we're seeing today is a result of that.

### **Operator**

Our next question comes from Brent Thill from Jefferies.

#### **Brent John Thill**

*Jefferies LLC, Research Division*

David, on the 25% to 30% sales capacity that you're seeing now, is that what you're continuing to plan for, for the year? Do you expect that growth rate to fade as you get into the back half of the year? Just give us a sense of the shape of that quota-carrying ramp in 2025.

#### **David M. Obstler**

*Chief Financial Officer*

We're continuing doing that. But what we see is that the ramp capacity because we essentially have this 6-month, 12-month ramp. If our plans are realized, the investment in headcount will translate into a higher ramp capacity increase. That's dependent upon getting the right people, training and enabling the right people and retaining them. But that's the idea, and we're hopeful for that.

### **Operator**

Our next question comes from Kirk Materne from Evercore ISS.

#### **Kirk Materne**

*Evercore ISI Institutional Equities, Research Division*

Oli, can you just talk about whether you've seen any discernible change in terms of the AI natives that you're working with in terms of what you're helping them with, meaning I assume it's -- majority is training at this point in time. But I was curious if you're seeing inferencing come into the equation in a bigger way or even if some of your bigger enterprise customers are starting to think a little bit more about inferencing. Just trying to get a sense on where we stand on that front.

#### **Olivier Pomel**

*Co-Founder, CEO & Director*

For the AI natives, actually, what we help them with mostly is not training. It's running their applications and their inference workloads as customer-facing. Because what's training for the AI natives tends to be largely homegrown one-off and different from -- between each and every one of them.

We expect that if most other companies and enterprises do significant training, that this will not be the case. This will not be one-off and homegrown. But right now, it is still the AI natives that do most of the training, and they still do it in a way that's largely homegrown. So when we see growth on the AI-native cohorts, that's growth of AI adoption because that's growth of customer-facing workloads by and large.

#### **Kirk Materne**

*Evercore ISI Institutional Equities, Research Division*

Okay. And then just one follow-up, maybe this is for David. Just -- I realize you guys are in investment mode right now. I'm just kind of curious, internally, how do you think about AI from an efficiency perspective? Meaning, I think Oli addressed it a little bit on the R&D side, but we've seen from other software companies that they are starting to see real yield in terms of using AI internally that might be masked with you all given the -- just the growth rate on your investments right now. But I'm just kind of curious how you think about that perhaps longer term relative to where your margins are today versus where they could be over a longer period of time.

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

Yes. And for right now, I think we're seeing the returns in productivity, whether that be salespeople getting more information or R&D. We're essentially trying to create an environment where we're encouraging the various departments to use it and learning from it. Long term, there might well be efficiency gains -- there may be efficiency gains that can be manifested in headcount. But right now, we're essentially trying to make the headcount investment that we are making more efficient and more...

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

Remember, we're constrained by sales capacity, and we're constrained by engineering capacity. So any gain of productivity we gain right now are basically reinvested in producing more, whether that's on the revenue side or on the product and engineering side.

Long term, obviously, we are strong believers in the fact that AI is a big driver of efficiency. And we can see how it can be -- if you squint, you see that it can be 10x cheaper to produce the same amount of software and things like that. So I think we definitely are headed there.

**Operator**

Our next question comes from Howard Ma from Guggenheim.

**Howard Ma**  
*Guggenheim Securities, LLC, Research Division*

I was hoping you guys could reconcile the consumption downtick in the enterprise segment versus the strong new logo ARR growth that I believe you said was over 70% year-over-year.

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

Yes. I mean there's no contradiction there. I think there's always variability in the usage. We see that month-to-month, we see that quarter-to-quarter. And what we see there is within the bounds of that.

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

Yes. We obviously -- we take a strong look. So I think the first one is what we talked about already, which is product sales capacity, winning market share in the market. The second one is we've always said -- we have some volatility now. Back when we had the post-COVID, we saw some correlation around broad parts of our customer base, indicating that were broader trends. In this case, we've looked at it and it's sort of volatility of usage, but not a broad trend.

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

Yes. And remember also, when we say we have strong bookings, it takes time for bookings to turn into revenue for us. For example, when we have a great new logo deal, we mentioned a few on the call, that usage will take many months to materialize and revenue will take months to materialize. But then, the back end of that, we expect to see it go far beyond what was the initial land with those customers. So there's definitely a timing effect there.

**Operator**

Our next question comes from Matt Hedberg from RBC.

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

I guess, customers can obviously change usage trends due to macro...

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

We can't hear you. I'm sorry, we can't hear you.

**Yuka Broderick**  
*Investor Relations*

Yes. You're breaking up.

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

Let me try this. Philosophically speaking, customers optimized in 2022 and 2023, are they still operating at fairly optimized levels?

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

Well, customers, I think, are running tighter optimization cycles now. We -- so we don't think -- the vast majority of customers don't build as much overhang of unoptimized usage that they need to come back to later. We do still see the same regular cycles, right? Typically, when they renew their cloud commitments or when they renew the observability commitment is when they're going to look at what they need to optimize and come back to it. But we think in general that, that cycle has become much tighter for customers.

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

Got it. And maybe just a quick one for David. It sounds like \$10 million of expenses for Eppo and Metaplane. Curious if your guidance includes any impact from a revenue perspective for those 2.

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

Well, for the second half of the year, we did not change the revenue guidance. So I would say, net-net, when you put everything together, we did not include it.

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

Yes. Impact is small.

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

And the impact is -- yes, and I think that's important to say. The impact is small, the combination of the revenues in the quarterly -- in the second half based on what we acquired, now we're obviously going to try to accelerate that. It is small anyway. Yes.

**Operator**

Our next question comes from Brad Reback from Stifel.

**Brad Robert Reback**  
*Stifel, Nicolaus & Company, Incorporated, Research Division*

Oli, there's been growing kind of customer focus on bring your own cloud. How should we think about your -- well, I should -- what are you thinking philosophically as it relates to that, maybe on-prem workloads in general?

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

Well, I think it's 2 different things. There is on-prem workloads, because some customers still have a lot of large on-prem workloads that are not going anywhere, and we're seeing that. And we have new product offerings that go squarely after that. So we can cover from the cloud, the on-prem workloads in a way that is cost effective for them and that they can use in conjunction with the monitoring of their own cloud environment. So that's one thing, and we are addressing that. We see a big opportunity there.

The second thing is the desire for customers to run on their own infrastructure, whether it's clouds or some on-premise version of a cloud. We think there's definitely a future where that becomes more important, especially if there's more geopolitical separation between the various parts of the world. Philosophically, we go wherever customers want us to go. So if that's the way they want to consume software in the end, that's definitely the way we're going to deliver it.

We are testing right now some ways for customers to manage part of their data directly on their own cloud. We have a new product that is being -- in customers' hands for that. And if we see strong market pull for that and if we see the world going in that direction, we'll definitely lean harder into that. So for us, this can be a big opportunity in the future.

**Operator**

In the interest of time, we have 1 more question. Our next question comes from Andrew Sherman from TD Cowen.

**Andrew Michael Sherman**  
*TD Cowen, Research Division*

The net new adds, 100,000 plus of 160 were the strongest since 4Q '22. Maybe just talk about the drivers of expansion there across new use cases, new products, share gains. And it sounds like some boomerang customers coming back to you as well.

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

Yes, all of the above. That's good. And look, I'm happy this number is high, but it's -- the reality is that it's consistent. And sometimes a bit lower, sometimes a bit higher, but all of our customers are going up the curve with us as they consolidate more and more of what they're doing into our platform and bring us more and more of their workloads, and they grow with us.

Of course, we're also very happy to see some customers coming back. So the one customer we mentioned on the call here has decided to sort of build their own using some commercial open source, and they decided to come back, and it's a motion we do see. We are very happy to see that in the end, from customers who've tried everything, including us in the competition, we're seen as being the best long-term choice. So -- and maybe we can close it on that.

**Operator**

Thank you. The question-and-answer session is now closed. I will now turn it over to Olivier Pomel for closing remarks.

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

All right. Thank you very much. So again, I want to thank all of our employees and customers for working with us and for us during the quarter. And I remind everyone that we are hard at work preparing our conference, DASH, on June 10 and 11 in New York. We have a lot of exciting stuff to show there, so we expect to see as many of you there as possible.

**Operator**

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

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

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

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