

Datadog, Inc. NasdaqGS:DDOG

FQ4 2025 Earnings Call Transcripts

Tuesday, February 10, 2026 8:00 AM ET

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

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

Peter Weed
*Bernstein Institutional Services LLC,
Research Division*

Brent John Thill
Jefferies LLC, Research Division

Raimo Lenschow
Barclays Bank PLC, Research Division

Gabriela Borges
*Goldman Sachs Group, Inc., Research
Division*

Sanjit Kumar Singh
Morgan Stanley, Research Division

Howard Ma
*Guggenheim Securities, LLC, Research
Division*

Todd Adair Coupland
*CIBC Capital Markets, Research
Division*

Ittai Kidron
*Oppenheimer & Co. Inc., Research
Division*

Koji Ikeda
BofA Securities, Research Division

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

Matthew George Hedberg
*RBC Capital Markets, Research
Division*

Presentation

Operator

Good day, and welcome to the Q4 2025 Datadog's earnings conference call. [Operator Instructions] As a reminder, this call may be recorded.

I would now like to turn the call over to Yuka Broderick, Senior Vice President of Investor Relations. Please go ahead.

Yuka Broderick *Investor Relations*

Thank you, Michelle. Good morning, and thank you for joining us to review Datadog's fourth 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 first quarter and fiscal year 2026 and related notes and assumptions, our product capabilities and our ability to capitalize on market opportunities. The words anticipate, believe, continue, estimate, expect, intend, will and similar expressions are intended to identify forward-looking statements or similar indications of future expectations.

These statements reflect our views today and are subject to a variety of risks and uncertainties that could cause actual results to differ materially. For a discussion of the material risks and other important factors that could affect our actual results, please refer to our Form 10-Q for the quarter ended September 30, 2025. Additional information will be made available in our upcoming Form 10-K for the fiscal year ended December 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 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 this morning to go over what was a very strong Q4 and overall, a really productive 2025. Let me begin with this quarter's business drivers. We continue to see broad-based positive trends in the demand environment. With the ongoing momentum of cloud migration and digital transformation. We experienced strength across our business, across our product lines and across our diverse customer base. We saw a continued acceleration of our revenue growth. This acceleration was driven in large part by the inflection of our broad-based business outside of the AI-native group of customers we discussed in the past. And we also continue to see very high growth within this AI-native customer group as they go into production and grow in users, tokens and new products.

Our go-to-market teams executed to a record \$1.63 billion in bookings, up 37% year-over-year. This included some of the largest deals we have ever made. We signed 18 deals over \$10 million in TCV this quarter, of which 2 were over \$100 million and 1 was an 8-figure land with a leading AI model company.

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 Q4 financial performance and key metrics, revenue was \$953 million, an increase of 29% year-over-year and above the high end of our guidance range. We ended Q4 with about 32,700 customers, up from about 30,000 a year ago. We also ended Q4 with about 4,310 customers with an ARR of \$100,000 or more, up from about 3,610 a year ago. These customers generated about 90% of our ARR. And we generated free cash flow of \$291 million with a free cash flow margin of 31%.

Turning to product adoption. Our platform strategy continues to resonate in the market. At the end of Q4, 84% of customers use 2 or more products, up from 83% a year ago. 55% of customers use 4 or more products, up from 50% a year ago, 33% of our customers use 6 or more products, up from 26% a year ago, 18% of our customers use 8 or more products, up from 12% a year ago, and as a sign of continued penetration of our platform, 9% of our customers use 10 or more products, up from 6% a year ago.

During 2025, we continued to land and expand with larger customers. As of December 2025, 48% of the Fortune 500 are Datadog customers. We think many of the largest enterprises are still very early in their journey to the cloud. The median Datadog ARR for our Fortune 500 customers is still less than \$0.5 million, which leaves a very large opportunity for us to grow with these customers. So we're landing more customers and delivering more value, and we also see that with the ARR milestones we're reaching with our products.

We continue to see strong growth dynamics with our core 3 pillars of observability: infrastructure monitoring, APM and log management as customers are adopting the cloud, AI and modern technologies. Today, infrastructure monitoring contributes over \$1.6 billion in ARR. This includes innovations to deliver visibility and insights across our customers' environments, whether they are on-prem, virtualized servers, containerized hosts, serverless deployments or parallelized GPU fleets.

Meanwhile, Log Management is now over \$1 billion in ARR. And this includes continued rapid growth with Flex Logs, which is nearing \$100 million in ARR. And our third pillar, the end-to-end suite of APM and DEM products also crossed \$1 billion in ARR. This includes an acceleration of our core APM product into the mid-30% year-over-year and currently our fastest-growing core pillar. We have now enabled our customers with the easiest onboarding and implementation in the market, while delivering unified deep end-to-end visibilities into their applications.

Now remember that even with these 3 pillars, we are still just getting started as about half of our customers do not buy all 3 pillars from us, or at least not yet.

Moving on to R&D and what we built in 2025. We released over 400 new features and capabilities this year. That's too much for us to cover today, but let's go over just some of our innovations. We are executing relentlessly on our very ambitious AI road map, and I will split our AI efforts into 2 buckets: AI for Datadog and Datadog for AI. So first, let's look at AI for Datadog. These are AI products and capabilities that make the Datadog platform better and more useful for customers. We launched Bits AI SRE Agent for general availability in December to accelerate root cause analysis and incident response. Over 2,000 trial and paying customers have run investigations in the past month, which indicates significant interest and shows great outcomes with Bits AI SRE. And we're well on our way with Bits AI Dev agent, which detects code level issues, generates fixes in production context and can even help release and monitor a fix and Bits AI Security Agent, which autonomously triages SIEM signals, conducts investigations and delivers recommendations.

The Datadog MCP server is being used by thousands of customers in preview. Our MCP server responds to AI agent and user prompts and uses real-time production data and rich Datadog context to drive troubleshooting, root cause analysis and automation. And we're seeing explosive growth in MCP usage, with the number of tool calls growing 11-fold in Q4 compared to Q3.

Second, let's talk about Datadog for AI. This includes capabilities that deliver end-to-end observability and security across the AI stack. We are seeing an acceleration in growth for LLM Observability. Over 1,000 customers are using the product and the number of spans since has increased 10x over the last 6 months. In 2025, we broadened the product to better support rapid agent development and iteration adding capabilities such as LLM Experiments and LLM Playground, LLM prompt analysis and custom LLM-as-a-judge. And we will soon release our AI Agents console to monitor usage and adoption of AI agents and coding assistants.

We are working with design partners on GPU monitoring, and we are seeing GPU usage increase in our customer base overall. And we are building into our products the ability to secure the AI stack against prompt injection attacks, model hijacking and data poisoning among many other risks.

Overall, we continue to see increased interest among our customers in next-gen AI. Today, about 5,500 customers use one or more Datadog AI integrations to send us data about their machine learning, AI and LLM usage.

In 2025, our Observability platform delivered deeper and broader capabilities for our customers. We reached a major milestone of more than 1,000 integrations, making it easy for our customers to bring in every type of data they need and engage with the latest technologies from cloud to AI. In log management, we're seeing success with our consolidation motion. During 2025, we saw an increasing demand to replace a large legacy vendor with takeouts in nearly 100 deals for tens of millions of dollars of new revenue. And we improved log management with notebooks, reference tables, log patterns, calculated fields and an improved Live Tail among many other innovations.

We launched Data Observability for general availability. Data is becoming even more critical in the AI era. With Data Observability, we are enabling end-to-end visibility across the entire data life cycle. We launched Storage Management last month, providing granular insights into cloud storage and recommendations to reduce spend. We delivered Kubernetes Autoscaling, so users can quickly identify overprovisioned clusters and deployments and right-size their infrastructure.

In the digital experience monitoring area, we launched Product Analytics to help product designers make better design decisions with clear data about user experience and behavior. And we delivered RUM without Limits, giving front-end teams full visibility into user traffic and performance and dynamically choosing the most useful sessions to retain.

In security, we are seeing increasing traction and are actively displacing existing market-leading solutions with Cloud SIEM in large enterprises. This year, our engineers shipped many new capabilities, including a tripling of the amount of content packs built into the product, and most importantly, the tight integration with Bits AI Security Agent, which has already shown promise as a strong differentiator in the market. We launched Code Security, enabling customers to detect and remediate vulnerabilities in their code and open source libraries from development to production. And we continue to advance our cloud security offering, adding Infrastructure as Code or IaC security, which detects and resolves security issues with Terraform. And we launched our security graph to identify and evaluate attack paths.

In software delivery, in January, we launched Feature Flags. They combine with our real-time observability to enable canary rollouts, so teams can deploy new code with confidence. And we expect them to gain importance in the future as they serve as a foundation for automating the validation and release of applications in an AI agentic development world.

We are also building out our internal developer portal, which includes software catalog and scorecards to help developers navigate infrastructure and application complexity, provide rich context to AI development agents and ultimately enable a faster release cadence. In Cloud Service Management, we launched OnCall, and now support over 3,000 customers with their incident response processes. And I already mentioned Bits AI SRE Agent, which pairs with OnCall to accelerate our customer incident resolution.

As you can tell, we've been very busy, and I want to thank our engineers for a very productive 2025. And most importantly, I'm even more excited about our plans for 2026. So let's move on to sales and marketing.

I want to highlight some of the great deals we closed this quarter. First, we landed an 8-figure annualized deal and our biggest new logo deal to date with one of the largest AI foundational model companies. This customer has a fragmented observability stack and cumbersome monitoring workflows leading to poor productivity. This is a consolidation of more than 5 open source, commercial, hyperscaler and in-house observability tools into the unified Datadog platform that has returned meaningful time to developers and has enabled a more cohesive approach to observability. This customer is experiencing very rapid growth. Datadog allows them to focus on product development and supporting their users, which is critical to their business success.

Next, we welcomed back a European data company in a nearly 7-figure annualized deal. This customer's log-focused observability solution had poor user experience and integrations, which led to limited user adoption and gaps in coverage. By returning to Datadog and consolidating 7 observability tools, they expect to reduce tooling overheads and improve engineering productivity with faster incident resolution. They will adopt 9 Datadog products at the start, including some of our newer products such as Flex Logs, Observability Pipelines, Cloud Cost Management, Data Observability and OnCall.

Next, we signed an 8-figure annualized expansion with a leading e-commerce and digital payments platform. This customer's products have an enormous reach and it's commercial APM solution had scaling issues, lacked correlation across silos and had a pricing model that was difficult to understand or predict. With this expansion, they are standardizing on Datadog APM using OpenTelemetry so their teams can correlate metrics, traces and logs to detect and resolve issues faster. And they have already seen meaningful impact with a 40% reduction in resolution times by their own estimates. This customer has adopted 17 products across the Datadog platform.

Next, we signed a 7-figure annualized expansion for an 8-figure annualized deal with a Fortune 500 food and beverage retailer. This long-time customer uses the Datadog platform across many products, but still has over 30 other observability tools and embarked on consolidating for cost savings and better outcomes. With this expansion, Datadog log management and Flex Logs will replace their legacy logging product for all ops use cases, with expected annual savings in the millions of dollars. This customer is expanding to 17 Datadog products.

Next, we signed a 7-figure annualized expansion with a leading health care technology company. This company was facing reliability issues, impacting clinicians during critical workflows and putting customer trust at risk. The customer will consolidate 6 tools and adopt 7 Datadog products, including LLM Observability to support their AI initiatives as well as Bits AI SRE Agent to further accelerate incident response.

Next, we signed an 8-figure annualized expansion, more than quadrupling the annualized commitment with a major Latin American financial services company. Given its successful tool consolidation projects and rapid adoption of Datadog products across all of its teams, this customer renewed early with us while expanding to additional products, including Data Observability, CI Visibility, Database Monitoring and Observability Pipelines. With Datadog, this customer showed measurable improvements in cost, efficiency,

customer experience and conversion rates across multiple lines of business. That proof of value led them to broaden their commitment with us, and has firmly established Datadog as their mission-critical observability partner.

Last and not least, we signed a 7-figure annualized expansion for an 8-figure annualized deal with a leading fintech company. With this expansion, the customer is moving their log data on to our unified platform. So teams can correlate telemetry in one place and save between hours and weeks in time to resolution for incidents. This customer has adopted 19 Datadog products across the platform, including all 3 pillars as well as digital experience, security, software delivery and service management. And that's it for our wins. Congratulations to our entire go-to-market organization for a great 2025 and a record Q4.

It was inspiring to see the whole team at our Sales Kick off last month and really exciting to embark on a very ambitious 2026.

Before I turn it over to David for a financial review, I want to say a few words on our longer-term outlook. There is no change to our overall view that digital transformation and cloud migration are long-term secular growth drivers for our business. So we continue to extend our platform to solve our customers' problems from end to end across their software development, production, data stack, user experience and security needs. Meanwhile, we're moving fast in AI, by integrating AI into the Datadog platform to improve customer value and outcomes and by building products to observe, secure and act across our customers' AI stacks.

In 2025, we executed very well and delivered for our customers against their most complex mission-critical problems. Our strong financial performance is an output of that effort. And we're even more excited about 2026 as we are starting to see an inflection in AI usage by our customers into their applications and as our customers begin to adopt our AI innovations, such as our Bits AI SRE Agent.

To hear about all that in detail and much more, I welcome you all to join us at our next Investor Day this Thursday in New York between 1:00 and 5:00 p.m.

I'll be joined by our product and go-to-market leaders to share how we are serving our customers, how we innovate to broaden our platform and how we are delivering greater value with AI. For more details, please refer to the press release announcing the event or head to investors.datadoghq.com. And with that, I will turn it over to our CFO, David.

David M. Obstler
Chief Financial Officer

Thanks, Olivier. Our Q4 revenue was \$953 million, up 29% year-over-year and up 8% quarter-over-quarter. Now to dive into some of the drivers of our Q4 revenue growth. First, overall, we saw robust sequential usage growth from existing customers in Q4. Revenue growth accelerated with our broad base of customers, excluding the AI natives to 23% year-over-year, up from 20% in Q3.

We saw strong growth across our customer base with broad-based strength across customer size, spending bands and industries. And we have seen this trend of accelerated revenue growth continue in January. Meanwhile, we are seeing continued strong adoption amongst AI-native customers with growth that significantly outpaces the rest of the business. We see more AI-native customers using Datadog with about 650 customers in this group. And we are seeing these customers grow with us, including 19 customers spending \$1 million or more annually with Datadog.

Among our AI customers are the largest companies in this space, as today 14 of the top 20 AI-native companies are Datadog customers.

Next, we also saw continued strength from new customer contribution. Our new logo bookings were very strong again this quarter, and our go-to-market teams converted a record number of new logos and average new logo land sizes continues to grow strongly.

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

And now moving on to our financial results. First, billings were \$1.21 billion, up 34% year-over-year. Remaining performance obligations, or RPO, was \$3.46 billion, up 52% year-over-year. And current RPO growth was about 40% year-over-year. RPO duration increased year-over-year as the mix of multiyear deals increased in Q4. We continue to believe revenue is a better indication of our business trends than billing and RPO.

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

First, our Q4 gross profit was \$776 million with a gross margin percentage of 81.4%. This compares to a gross margin of 81.2% last quarter and 81.7% in the year ago quarter. Q4 OpEx grew 29% year-over-year versus 32% last quarter and 30% in the year ago quarter. And we continue to grow our investments to pursue our long-term growth opportunities, and this OpEx growth is an indication of our successful execution on our hiring plans. Our Q4 operating income was \$230 million for a 24% operating margin compared to 23% last quarter and 24% in the year ago quarter.

Now turning to the balance sheet and cash flow statements. We ended the quarter with \$4.47 billion in cash, cash equivalents and marketable securities. Cash flow from operations was \$327 million in the quarter. After taking into consideration capital expenditures and capitalized software, free cash flow was \$291 million for a free cash flow margin of 31%.

And now for our outlook for the first quarter and the full fiscal year 2026. 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.

For the first quarter, we expect revenues to be in the range of \$951 million to \$961 million, which represents a 25% to 26% year-over-year growth. Non-GAAP operating income is expected to be in the range of \$195 million to \$205 million, which implies an operating margin of 21%. Non-GAAP net income per share is expected to be in the \$0.49 to \$0.51 per share range based on approximately 367 million weighted average diluted shares outstanding. And for the full fiscal year 2026, we expect revenues to be in the range of \$4.06 billion to \$4.10 billion which represents 18% to 20% year-over-year growth. This includes modeling within our guidance that our business, excluding our largest customer, grows at least 20% during the year.

Non-GAAP operating income is expected to be in the range of \$840 million to \$880 million, which implies an operating margin of 21%. And non-GAAP net income per share is expected to be in the range of \$2.08 to \$2.16 per share based on approximately 372 million weighted average diluted shares.

Finally, some additional notes on our guidance. First, we expect net interest and other income for the fiscal year 2026 to be approximately \$140 million. Next, we expect cash taxes in 2026 to be about \$30 million to \$40 million and we continue to apply a 21% non-GAAP tax rate for 2026 and beyond.

Finally, we expect capital expenditures and capitalized software together to be in the 4% to 5% of revenue range in fiscal year 2026.

To summarize, we are pleased with our strong execution in 2025. Thank you to the Datadog teams worldwide for a great 2025, and I'm very excited about our plans for 2026.

And finally, we look forward to seeing many of you on Thursday for our Investor Day. And now with that, we will open up our call for questions. Operator, let's begin the Q&A.

Question and Answer

Operator

[Operator Instructions] Our first question comes from Sanjit Singh with Morgan Stanley.

Sanjit Kumar Singh

Morgan Stanley, Research Division

Congrats on a strong close of the year and a successful 2025. Olivier, I wanted to get your updated views in terms of where observability is headed. In the context of a lot of advancements when it comes to agentic frameworks, agentic deployments, the stuff that we've seen from Anthropic and new frontier models from OpenAI, just in terms of like what this means for observability as a category, defensibility of it in terms of can customers use these tools to build homegrown solutions for observability? So just give your latest comments on defensibility of the category, and how Datadog may potentially have to evolve in this new sort of a agentic era?

Olivier Pomel

Co-Founder, CEO & Director

Yes. I mean, look, the -- there's a few different ways to look at it. One is there's going to be many more applications than there were before. Like people are building much more and they are building much faster. We covered that in previous calls, but we think that the -- this is nothing, but an acceleration of the increase of productivity for developers in general, so you can build a lot faster. As a result, you create a lot more complexity because you build more than you can understand at any point in time. And you move a lot of the value from the act of writing the code, which now you actually don't do yourself anymore to validating, testing, making sure it works in production, making sure it's safe, making sure it interacts well with the rest of the world, with end users, make sure it does what it's supposed to do for the business, which is what we do with observability. So, we see a lot more volume there, and we see that as what we do basically where observability can help.

The other part that's interesting is that we -- a lot happens -- a lot more happens within these agents and these applications. And a lot of what we do as humans now starts to look like observability. Basically, we're here to understand -- we're trying to understand what the machine does. We're trying to make sure it's aligned with us. We're trying to make sure the output is what we expected when we started, and that we didn't break anything. And so, we think it's going to bring observability more widely in domains that it didn't necessarily cover before.

So we think that these are accelerants, and we -- I mean, obviously, we have a [horse] in this one, but we think that observability and the contact between the code, the applications and the real world and product environment and real users and the real business is the most interesting, the most important part of the whole AI development life cycle today.

Sanjit Kumar Singh

Morgan Stanley, Research Division

And maybe just one follow-up on that line of thinking. In a world where there's a greater mix between human SREs and agentic SREs, is there any sort of evolution that we need to think about in terms of whether it's UI or how workflows work in observability and how maybe Datadog sort of tries to align with that evolution that's likely to come in the next couple of years?

Olivier Pomel

Co-Founder, CEO & Director

Yes, there's going to be an evolution, that's certain. There's going to be a lot more automation. We see it today, like we see the -- all the signs we see point to everything moving faster, more data and more interactions, more systems, more releases, more breakage, more resolutions of those breakages, more bugs, more vulnerabilities, everything. So, we see an acceleration there. At the end of the day, humans will still have some form of UI to interact with all that. And a lot of the interaction will be automated by agent. So we're building the products to satisfy both conditions.

So, we have a lot of UIs, and we are able to present humans with UIs that represent how the world works, what their options are, give them familiar ways to go through problems and to model the world. And we also are exposing a lot of our functionality to agents directly. We mentioned on the call, we have an MCP server that is currently in preview and that is really seeing explosive growth of usage from our customers. And so, it's a very likely future that part of our functionality is delivered to agents through MCP servers or

the likes. Part of our functionality is directly implemented by our own agents, and part of our functionality is delivered to humans with UIs.

Operator

Our next question comes from Raimo Lenschow with Barclays.

Raimo Lenschow

Barclays Bank PLC, Research Division

Congrats from me as well. Staying on a little bit on that AI theme, Olivier, the 8-figure deal for a model company is really exciting. I assume they try to do it with some open source tooling, et cetera, but -- and actually went from like almost paying not a lot of money to paying you more money. What drove that thinking? What do you think what they saw that kind of convinced them to do that? And it's now the second one after the other very big model provider. So clearly, that whole debate in the market between, oh, you can do that on the cheap somewhere is not kind of quite valid. Could you speak to that, please?

Olivier Pomel

Co-Founder, CEO & Director

I mean the situation is just very similar to every single customer we land. Every customer we land has some -- has had some at homegrown. They have some open source. They might still run some open source, like that's typically where we see everywhere. It's cheaper to do it yourself is usually not the case. So, your engineers typically are very well compensated in the big part of the spend in this company. Their Velocity is what gates just about anything else in the business. And so usually, when we come in, when customers start engaging with us, we can very quickly show value that way. So, it's not any different from what we see with any other customer. And also within the AI cohort, it's not original at all like -- AI cohort in general is who's who of the companies that are growing very fast and that are shaping the world in AI and they're all adopting our product for the same reasons, sometimes the different volumes because those companies have different scales, but the logic is the same.

Operator

Our next question comes from Gabriela Borges with Goldman Sachs.

Gabriela Borges

Goldman Sachs Group, Inc., Research Division

Congratulations on the quarter. Oli, I wanted to follow up on Sanjit's question on how to think about where the line is between what an LLM can do longer term and the domain experience that you have in observability? If I think about some of Anthropic's recent announcements, they're talking about LLMs as a broader anomaly detection type tool, for example, on the security vulnerability management side. How do you think about the limiting factor to using LLMs as a n anomaly detection tool that could potentially take share from observability over time in the category? And how do you think about the moat that Datadog has that offers customers a better solution relative to where the road map in LLMs can go long term?

Olivier Pomel

Co-Founder, CEO & Director

Yes. So that's a very good question. We definitely see that LLMs are getting better and better, and we'll bet on them getting significantly better every few months as we've seen over the past couple of years. And as a result, they are very, very good at looking at broad sets of data. So, if you feed a lot of data to an LLM and ask for an analysis, you're very likely to get something that is very good and that is going to get even better. So, when you think of what we have that is fundamentally our moat here, there's 2 parts. One is how we are able to assemble that contact, so we can feed it into those intelligence engines. And that's how we aggregate all the data we get, we parse out the benefits. We understand how everything fits together and we can feed that into the LMM. That's in part what we do, for example, today, we expose these kinds of functionality behind our MCP server. And so, customers can recombine that in different ways using different intelligence tools. But the other part that we think where the world is going for observability is that right now, we are -- the [SDLC] is accelerating a lot, but it's still somewhat slow. And so it's okay to have incidents and run post-hoc analysis on those incidents and maybe use some outside tooling for them.

Where the world is going is you're going to have many more changes, many more things. You cannot actually afford to have incidents to look at for everything that's happening in your system. So, you need to be proactive. You'll need to run analysis in stream as all the data flows through, you'll need to run detection and resolution before you actually have outages materialize. And for that, you'll need to be embedded into the data plane, which is what we run. And you also need to be able to run specialized models that can act on that

data as opposed to just taking everything and summarizing everything after the fact 10, 15 minutes later. And that's what we're uniquely positioned to do.

We are building that. We're not quite there yet, but we think that a few years from now, that's what the world is going to run, and that's what makes us significantly different in terms of how we can apply anomaly detection, intelligence and preemptive resolution into our systems.

Gabriela Borges
Goldman Sachs Group, Inc., Research Division

That makes a lot of sense. My follow-up...

Olivier Pomel
Co-Founder, CEO & Director

By the way, the data points we're talking about are very real time, and there are many others of magnitude larger in terms of data flows, data volumes than what you typically feed into an LLM. So it's a bit of a different problem to solve.

Gabriela Borges
Goldman Sachs Group, Inc., Research Division

Yes. Super interesting. My follow-up for both you, Oli and David, you've mentioned a couple of times now some of the conversations you have with customers about value creation within the Datadog platform. Tell us a little bit about how some of those conversations evolve when the customer sees that in order to do observability for more AI usage, perhaps that Datadog bill is going up. What are some of the steps that you can take to make sure the customer still feels like they're getting a ton of value out of the Datadog platform?

Olivier Pomel
Co-Founder, CEO & Director

Well, there's a few things. I mean, first, again, the rule of software always applies. There's only 2 reasons people buy your product is to make more money or to save money. So, whatever you do, when customers use a new product, they need to see a cost savings somewhere or they need to see that they're going to get to customers they wouldn't get to otherwise. So, we have to prove that. We always prove that. Any time a customer buys a product, that's what is happening behind the scenes. The -- in general, when customers add to our platform as opposed to bringing another vendor in or another product in, they also spend less by doing it on our platform.

Operator

Our next question comes from Ittai Kidron with Oppenheimer & Company.

Ittai Kidron
Oppenheimer & Co. Inc., Research Division

Congrats, quite an impressive finish for the year. David, I wanted to dig in a little bit into your '26 guide. I just want to make sure I understand some of your assumptions. So maybe you could talk about the level of conservatism that you've built into the guide for the year? And also, you've talked about at least 20% growth for the core, excluding the largest customer, but what is it that we should assume for the large customer? And now when you look at the AI cohort, excluding this large customer, are there any concentrations evolving over there given your strong success there?

David M. Obstler
Chief Financial Officer

Yes. There are 3 questions in there. The first is overall on guidance, except what we're going to speak about next, we took the same approach as we looked at the organic growth rates and the attach rates and then the logo accumulation rates and discounted that. So for the overall business, which is quite diversified, we talked about diversification by industry, by geography, by SMB, mid-market and enterprise, we took the same approach. We noted that with the guidance being 18% to 20% and the non-AI or heavily diversified business being 20% plus, that would imply that the growth rate of that core business assumed in the guidance is higher than the growth rate of the large customer. It doesn't mean the large customer is growing any which way. It's just that in our consumption model, we essentially don't control that. And so, we made a very conservative assumption there.

And the last point, I think you mentioned is the highly diversified. We said 650 names in the AI cohort is quite diversified, essentially would be very similar to our overall business, which we have a range of customers, but not the concentration level. And what we're

seeing there is significant growth. But like our overall distributed customer base, a growth and then potentially some working on how the product is being used, but nothing out of the ordinary relative to the overall customer base in the very diversified AI set of customers outside the largest customer. Hopefully that's helpful.

Ittai Kidron

Oppenheimer & Co. Inc., Research Division

Okay. That's great. Yes. And can you give us the percent of revenue of the AI cohort this quarter?

David M. Obstler

Chief Financial Officer

We didn't -- have not put it in there.

Operator

Our next question comes from Todd Coupland with CIBC.

Todd Adair Coupland

CIBC Capital Markets, Research Division

I wanted to ask you about competition and how the LLM rise is impacting share shifts. Just talk about that and how Datadog will be impacted?

Olivier Pomel

Co-Founder, CEO & Director

Yes. I mean, there hasn't been -- in the market with customers, there hasn't been any particular change in competition in that we see the same kind of folks and the positions are relatively similar. And we are pulling away. We're taking share from anybody who has scale. And I know there's been noise. There were a couple of M&A deals that came up, and we got some questions about that. The companies in there were not particularly winning companies, nothing that we saw in deals, nothing that had a large market impact. And so we don't see that as changing the competitive dynamics for us in the near future.

We also know that competing in observability is a very, very full-time job. It's a very innovative market. And we know exactly what it is we have to do and have to do to keep pulling away the way we are. And so we're very confident in our approach and we're going to do in the future there.

With the rise of LLM, there's clearly more functionality to build and there are new ways to serve customers. We mentioned our LLM Observability product. There are a few other products in the market for that. I think it's still very early for that part of the market, and that market is still relatively undifferentiated in terms of the kinds of products they are, but we expect that to shake out more into the future. We think, in the end, there's no reason to have observability for your LLM that is different from the rest of your system in great part because your LLM don't work in isolation. The way they implement their smarts is by using tools, the tools on your applications and your existing applications or new applications you build for that purpose. And so you need everything to be integrated in production, and we think we stand on a very strong footing there.

Operator

Our next question comes from Mark Murphy with JPMorgan.

Mark Ronald Murphy

JPMorgan Chase & Co, Research Division

Olivier, Amazon is targeting \$200 billion in CapEx this year. If you include Microsoft and Google, that CapEx is going to exceed \$500 billion this year for the big 3 hyperscalers, it's growing 40% to 60%. I'm wondering if you've collected enough signal from the last couple of years of CapEx, that trend to estimate how much of that is training related and when it might convert to inferencing where Datadog might be required? In other words, are you looking at this wave of CapEx and able to say it's going to create a predictable ramp in your LLM observability revenue? Maybe what inning of that are we in? And then I have a follow-up.

Olivier Pomel

Co-Founder, CEO & Director

I think it's pretty too reductive to peg that on LLM observability. I think it points to way more applications, way more intelligence, way more of everything into the future. Now it's kind of hard to directly map the CapEx from those companies into what part of the infrastructure is actually going to be used to deliver value 2 or 3 or 4 years from now. So I think we'll have to see what the conversion rate is on that. But look, it definitely points to very, very, very large increases in the complexity of the systems, the number of systems and the reach of the systems in the economy. And so we think it's going to be -- like it's going to be of great help to our business, let's put it this way.

Mark Ronald Murphy
JPMorgan Chase & Co, Research Division

Yes. Great help. Okay. And then as a quick follow-up, there is an expectation developing that OpenAI is going to have a very strong competitor, which is Anthropic kind of closing the gap, producing nearly as much revenue as OpenAI in the next 1 to 2 years. You mentioned an 8-figure land with an AI model company. I'm wondering, if we step back, do you see an opportunity to diversify that AI customer concentration, whether sometimes it might be a direct customer relationship there? Or it could be some of the products like Claude Code being adopted globally, just kind of creating more surface area to drive business to Datadog. Can you comment on maybe what is happening there among the larger AI providers or whether you can diversify that out?

Olivier Pomel
Co-Founder, CEO & Director

Yes. I mean, look, we've never been -- we're not built as a business to be concentrated on a couple of customers. That's not how we've become successful. That's probably not how we'll be successful in the long term. So yes, I mean we -- at the end of the day, it should be irrational for customers -- for all customers in the AI cohort not to use our product. So we see -- we have some great successes with the customers currently in that cohort. We see more. By the way, we have more that are more inbound there and more customers that are talking to us from the largest even hyperscaler level AI labs. And we expect to drive more business there in the future. I think there's no question about that.

David M. Obstler
Chief Financial Officer

And you're seeing that in some of the metrics we've been giving in terms of the number of AI-native customers, the size of some of these customers. So to echo what Oli said, we are essentially selling to many of the largest players, which results in greater size of the cohort and more diversification.

Operator

Our next question comes from Matt Hedberg with RBC.

Matthew George Hedberg
RBC Capital Markets, Research Division

Congrats from me as well. David, a question for you. Your prior investments are clearly paying off with another quarter of acceleration, and it seems like you're going to continue to invest in front of the future opportunity. I think op margins are down maybe 100 basis points on your initial guide. I'm curious if you can comment on gross margin expectations this year, and how you also might realize incremental OpEx synergies by using even more AI internally?

David M. Obstler
Chief Financial Officer

Yes. On the gross margin, I think what we said is plus or minus the 80% mark. We try to engineer when we see opportunities for efficiency, we've been quite good at being able to harvest them. At the same time, we want to make sure we're investing in the platform. So, I think what we're essentially -- where we are today is very much sort of in line with what we said we're targeting. There may be opportunities longer term, but we also are trying to balance those opportunities with investment in the platform.

And in terms of AI, to date, we are using it in our internal operations. So far, it's -- with the first signs of what we're seeing is productivity and adoption. We will continue to update everybody as we see opportunities in terms of the cost structure. Oli, anything else you want to go over?

Olivier Pomel
Co-Founder, CEO & Director

Yes. I mean, look, we -- the expectation in the short-, mid-term anyway should be that we keep investing heavily in R&D. We're getting a lot -- we see great productivity gains with AI there, but at this point of detail, it helps us build more faster and get to solve more problems for our customers. And -- but we're very busy adopting AI across the organization.

Operator

Our next question comes from Koji Ikeda with Bank of America.

Koji Ikeda

BofA Securities, Research Division

Olivier, maybe a question for you. A year ago, you talked about how -- while some customers do want to take observability in-house, it's really a cultural choice. It may not be rational unless you have tremendous scale, access to talent and growth is not limited by innovation bandwidth, which most companies do not. And so, it is a year later, and it does seem like the industry and the ecosystem and everything has changed quite a bit. So, I was hoping to get your updated views on these thoughts, if it has changed at all over the past year and why?

Olivier Pomel

Co-Founder, CEO & Director

No. I mean, look, it's something that happens sometimes, but it's a small minority of the cases. Like the general notion is customers start with some homegrown or attempts to do things themselves, then they move to our product, then they scale with our products. Sometimes they optimize a little bit along the way, but the general notion is they do more and more with us. They rely on us for more of their -- solving more of their problems, and they outsource the problem, and increasingly, the outcomes, to us. So, I don't think that's changing.

Look, we'll still see customers here and there that choose to in-source it and do it themselves, again, hugely for cultural reasons. I would say, economically or from a focused perspective, it doesn't make sense for the vast majority of companies. And we even see teams at hyperscalers that have all the tooling in the world, all the money in the world, all the know-how in the world and that still choose to use our product because it gives them a more direct path to solving their problems.

Operator

And our next question comes from Peter Weed with Bernstein Research.

Our next question comes from Brad Reback with Stifel.

Brad Robert Reback

Stifel, Nicolaus & Company, Incorporated, Research Division

Oli, the sustained acceleration in the core business is pretty impressive. Obviously, you all have invested very aggressively in go-to-market over the last kind of 18 to 24 months. Can you give us a sense of where you are in that productivity curve? And if there's additional meaningful gains, you think? Or is it incremental? And maybe where you see additional investments in the next 12 to 18 months?

Olivier Pomel

Co-Founder, CEO & Director

Yes. I mean we feel good about the productivity. I think the main drivers from us in the future is we still need to scale, and we're still scaling the go-to-market team. We're not at the scale we need to be in every single market segment we need to be in the world right now. And so we keep scaling there. So, the focus now is not necessarily to improve productivity, it's to scale while maintaining productivity. And of course, there's so many, many things we can do. Like we -- even though we love our performance, there's always a bunch of things that could be better, territories that could be better, productivity that could be better, things like that. So, we have tons of work -- tons of things we want to do, tons of things we want to fix, tons of things we want to improve. But overall, we feel good about what happened. We feel good about scaling, and you should expect more scaling for us on the go-to-market side in the year to come.

Operator

Our next question comes from Howard Ma with Guggenheim.

Howard Ma

Copyright © 2026 S&P Global Market Intelligence, a division of S&P Global Inc. All Rights reserved.

Guggenheim Securities, LLC, Research Division

I have one for Olivier. The core APM product, growing in the mid-30% growth. That is pretty impressive, and I think better than maybe a lot of us expected. The question is, is that a reacceleration? And is the growth driven by AI-native companies that are using Datadog's real user monitoring and other DEM features as compared to -- or as opposed to rather core enterprise customers that are building more applications?

Olivier Pomel
Co-Founder, CEO & Director

Yes. I think, I mean, look, APM, in general, I think has always been a bit of a steady eddy in terms of the growth, like it's a product that takes a little bit longer to deploy than others, which is further into the applications. And so, it takes a little bit longer to penetrate within the customer environment. That being said, we did -- a number of different things we did that helped with the growth there. One is we invested a lot in actually making that onboarding deployment a lot simpler and faster. So, we think we'll have the best in the market for that and it shows. Second, we invested a lot in the digital experience side of it. And it's very differentiated, something our customers love and is driving a lot of adoption of the broader APM suite, and we expect to see more of that in the future. And third, we make investments in go-to-market. We cover the market better. And so, we're getting into more looks at more deals in more parts of the world. And so, all of that combined helps that product reaccelerate growth quite a bit. And so, we feel actually very, very good about it, which is why we keep investing.

Overall, we still only have a small part of the pure APM market. That product is scaled at about \$10 billion, including DEM, but the market is larger. And so, we think there's a lot more we can do there.

David M. Obstler
Chief Financial Officer

I want to add, we talked about, as Oli just mentioned, that we're not penetrated across our customer base, and therefore, we're continuing to consolidate onto our platform. So, we have quite a number of wins where we already have other products. We already have infra and logs, and we're consolidating APM.

Howard Ma
Guggenheim Securities, LLC, Research Division

David, as a follow-up for you on margin, are the large AI-native customers significantly dilutive to gross margin? And when you think about the initial 2026 margin guide, how much of that reflects potentially lower gross margin tied to those customers versus incremental investments?

David M. Obstler
Chief Financial Officer

On a weighted average, they're not. As we've always said, for larger customers, it isn't about the AI-natives or non-AI-natives, it has to do with the size of the customer. We have a highly differentiated -- diversified customer base. So, I would say we're essentially expecting a similar type of discount structure in terms of size of customer as we have going forward. And there are consistent ongoing investments in our gross margin, including data centers and development of the platform. So, I think it's more or less what we've seen over the past couple of years, not really affected by AI or non-AI native.

Operator

Our next question comes from Peter Weed with Bernstein Research.

Peter Weed
Bernstein Institutional Services LLC, Research Division

Can you hear me this time?

Olivier Pomel
Co-Founder, CEO & Director

Yes.

Peter Weed
Bernstein Institutional Services LLC, Research Division

Yes, apologies for the last time. Great quarter. Looking forward, I think one of your most interesting exciting opportunities really is around Bits AI and I'd love to hear kind of like how you think that opportunity shapes up? Like how do you get paid the fair value for the productivity you're bringing to the SRE and the broader operations team and really how you see competition playing out in that space because obviously, we've seen start-ups coming in. There's questions about Anthropic and where they want to go. How does Datadog really capture this value and protect it for the business?

Olivier Pomel
Co-Founder, CEO & Director

Yes. I mean, look, the way we currently sell a lot of these products is you show like the difference in time spent. And when the alternative is you try and solve a problem yourself and you have an outage and you start a bridge and you have 20 people on the bridge and they look for 3 hours for the root cause, you know, and you wake up people in the middle of night for that. Like, it's very expensive. It takes a lot of time. There's a lot of customer impact because the outages are long. And if the alternative is in 5 minutes, you have the answer, you only get 3 people looking that are the right folks and you have a fix within 10 minutes, shorter impact on the customer, many, many, many less folks internally involved, lower cost. So, it's fairly easy to make that case. And so that's what we saw the value there.

Longer term, as I was saying earlier, I think the -- right now, the state-of-the-art for incident resolution is post-hoc. You know, you have an incident and you look into it. And you diagnose it and then you resolve it. So yes, maybe you could be the customer impact from 1 hour to 15 minutes. But you still have an issue, you still have impact, you still distract the team, you still have teams working on that. I think longer term, what's going to happen is the systems will get in front of issues. They will auto diagnose issues. They will help pre-mitigate or pre-remediate potential issues. And for that, the analysis will have to be run in stream, which is a very different thing. You can massage data and give it to an LLM for post-hoc analysis and a lot of the value is going to be in the gathering the data, but you also have quite a bit of value in the smarts that are done in the back end by the LLM for that. And that's something that is done by Anthropic, the opening eyes of the world today.

I think as you look at being in-stream looking at 3, 4, 5 orders of magnitude, more data, looking at the data in real time, and passing judgment in real time on what's normal, what's anomalous and what might be going wrong doing that hundreds, thousands, millions of times per second. I think that's what is going to be our advantage and where it's going to be much harder for others to compete, especially general purpose AI platforms.

Operator

Our next question comes from Brent Thill with Jefferies.

Brent John Thill
Jefferies LLC, Research Division

David, I think many gravitate back to that mid-20% margin you put up a couple of years ago. And I know the last couple of years, including the guide are looking at low 20%. Can you talk to maybe your true north, how you're thinking about that, obviously, growth being #1, but how you're thinking about the framework on the bottom line?

David M. Obstler
Chief Financial Officer

Yes, the framework is we try to plan with more conservative revenues, understanding that if the revenues exceed above the targets that we give, it's difficult in the short term to invest incrementally. So, what we're trying to do is invest first in the revenue growth and then layer in additional investment as we see -- if we see excess of target. So generally, it reflects, one, the continued investment, which we think is paying off, both in terms of the platform and R&D as well as in -- and including AI as a go-to-market. And then as we've seen over the years in our beat and raise, we've tended to have some of that flow through into the margin line and then re-up again for the next phase of growth.

Brent John Thill
Jefferies LLC, Research Division

And any big changes in the go-to-market or big investments you need to make, David, this year to address what's happened in the AI cohort or not?

David M. Obstler
Chief Financial Officer

We're continuing. It's very similar to what we're doing, which is to try to work with clients to prove value over time that reflects -- that manifests itself in our account management and our CSM as well as our enterprise. So no, I think for this year, we are looking at capacity growth, including geographic, deepening the ways we interact with customers, expanding channels, very much similar to what we've done in the previous years.

Olivier Pomel

Co-Founder, CEO & Director

That's going to be it for today. So on that, I'd like to thank all of you for listening to this call, and I think we'll meet many of you on Thursday for our Investor Day. So thank you all. Bye.

David M. Obstler

Chief Financial Officer

Thank you.

Operator

Thank you for your participation. You may now disconnect. Everyone, have a great day.

Copyright © 2026 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 (free of charge), and www.ratingsdirect.com and 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.
© 2026 S&P Global Market Intelligence.