

**Datadog, Inc. NasdaqGS:DDOG**  
**FQ3 2024 Earnings Call Transcripts**  
**Thursday, November 7, 2024 1:00 PM GMT**

# Table of Contents

Call Participants	.....	3
Presentation	.....	4
Question and Answer	.....	8

# Call Participants

## EXECUTIVES

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

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

**Yuka Broderick**  
*Investor Relations*

## ANALYSTS

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

**Gray Wilson Powell**  
*BTIG, LLC, Research Division*

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

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

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

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

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

# Presentation

## Operator

Good day, and thank you for standing by. Welcome to the Q3 2024 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, Senior Vice President of Investor Relations. Please go ahead.

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

Thank you, Liz. Good morning, and thank you for joining us to review Datadog's third quarter 2024 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 fourth quarter and fiscal year 2024 and related notes, our gross margins and operating margins, our product capabilities and our ability to capitalize on market opportunities. The words anticipate, believe, continue, estimate, expect, intend, will and similar expressions are intended to identify forward-looking statements or similar indications of future expectations. These statements reflect our views only as of today and are subject to a variety of risks and uncertainties that could cause actual results to differ materially.

For a discussion of the material risks and other important factors that could affect our actual results, please refer to our Form 10-Q for the quarter ended June 30, 2024. Additional information will be made available in our upcoming Form 10-Q for the fiscal quarter ended September 30, 2024, 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](https://investors.datadoghq.com).

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

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

Thanks, Yuka, and thank you all for joining us this morning. We are pleased to report on a solid Q3 as we continue to execute against our goals to help our customers grow faster, safer and more efficient as we modernize our applications. We kept broadening our platform in observability and beyond, including in next-gen AI, where interest continues to rise, and we added new customers while expanding with existing ones as they grow into the cloud.

Let me start with a review of our Q3 financial performance. Revenue was \$690 million, an increase of 26% year-over-year and above the high end of our guidance range. We ended the quarter with about 29,200 customers, up from about 26,800 a year ago. We had about 3,490 customers with ARR of \$100,000 or more, up from about 3,130 a year ago, and these customers generated about 88% of our ARR. And we generated free cash flow of \$204 million with a free cash flow margin of 30%.

Turning to platform adoption. Our platform strategy continues to resonate in the market. As of the end of Q3, 83% of customers were using 2 or more products, up from 82% a year ago. 49% of customers were using 4 or more products, up from 46% a year ago; 26% of our customers were using 6 or more products, up from 21% a year ago; and 12% of our customers were using 8 or more products, up from 8% a year ago.

We continue to execute on growth across the 3 pillars of observability. And we are pleased to report that infrastructure monitoring, our APM suite and log management together represent more than \$2.5 billion in ARR. As a reminder, within the APM suite, we include core APM, Synthetics, Real User Monitoring and Continuous Profiler.

We also want to call out our newer products, which are increasingly contributing to our business. Of our 23 products, 15 now exceeded \$10 million in ARR. These include each of our cloud security products as well as CI Visibility and Cloud Cost Management. So we have many products beginning to contribute to our revenue growth, and we're continuing to build greater capabilities within those products for our customers.

Now let's discuss this quarter's business drivers. Overall, the business environment for Datadog has remained stable and similar to what we have seen throughout 2024. Our customers overall are growing their cloud usage, while some are continuing to be cost conscious. In Q3, we continue to see existing customer usage growth broadly in line with our expectations. Our usage growth with

existing customers continued to be higher than in the year ago quarter, and we saw healthy growth across our product lines with newer products growing faster than more mature products off a smaller base.

Finally, churn continues to be low and gross revenue retention was stable in the mid- to high 90s, highlighting the mission-critical nature of our platform for our customers.

Moving on to R&D. In the next-gen AI space, customers continue to experiment with new AI technologies. And as they do, they want to get visibility into their AI use. At the end of Q3, about 3,000 customers used one or more Datadog AI integrations to send us data about their AI, machine learning and LLM usage. As some of these experiments start turning into production AI applications, we are seeing initial signs of traction for our LLM observability products.

Today, hundreds of customers are using LLM observability with more exploring it every day. And some of our first paying customers have told us that they have cut the time spent investigating LLM latency, errors and quality from days or hours to just minutes. Our customers not only want to understand the performance and cost of their LLM applications, they also want to understand the LLM model performance within the context of their entire application. So they are using APM alongside LLM observability to get fully integrated end-to-end visibility across all their applications and tech stacks.

Meanwhile, we continue to work to make the Datadog platform the best place for customers to monitor, secure and take action on their systems no matter where they deploy. In September, we launched Datadog monitoring for Oracle cloud infrastructure for general availability. With this launch, our customers gain visibility into their OCI stack and they can manage in real time the performance of OCI cloud services, servers, VMs, databases, containers and apps in Datadog. And customers can now unify their monitoring across OCI, other clouds and on-prem environments.

We also continue to extend our platform in new ways to bring value to our customers. At our DASH User Conference this summer, we announced Datadog OnCall, our newest product in the cloud service management space. As you know, our customers use Datadog extensively during their workdays for alerting and troubleshooting, whether that's for observability or security use cases. Now with Datadog OnCall, we are bringing a modern paging experience directly into our unified platform, and we now offer a completely integrated solution that covers incidents from end-to-end from detection, alerting and paging to incident management, troubleshooting and resolution.

Even though OnCall is still in limited availability, we are already seeing very strong reception for the product, and we are beginning to see customers request OnCall as part of their deals. In particular, new customers are interested in including paging as part of their land with Datadog.

So we're working hard to deepen and broaden our platform. And our innovations are rightfully being recognized by independent research firms. We are pleased to see that for the fourth year in a row, Datadog has been named a leader in the 2024 Gartner Magic Quadrant for Observability Platforms. We believe that this validates our approach to deliver a unified platform, which breaks down silos across teams. And Datadog has also been named a leader in Gartner's very first Magic Quadrant for Digital Experience Monitoring, which includes Datadog's products across synthetic testing, real user monitoring, product analytics, session replay and error tracking.

Now let's move on to sales and marketing. Our sales team continued to execute this quarter, and we added some exciting new customers while expanding with many more. So let's go through a few examples. First, we landed a 7-figure annualized deal with a leading e-commerce company in India. With its previous observability vendor, the customer saw quickly increasing costs while lacking the enterprise-grade observability they needed.

By switching to Datadog, they expect to support their scaling goals and will rely for that on our tracing, granular profiling and cloud integration support. I will note that we are pleased to have landed a large new logo customer in India, and we are continuing to invest to grow our presence and our opportunities there.

Next, we signed a 6-figure annualized land with a major U.S. Federal Agency. This agency is beginning to move some of its workloads to the cloud and is expanding the services offers to every single U.S. citizen through cloud applications. They have chosen Datadog to observe and secure their cloud environment and deliver a faster, better experience to end users. This deal includes 8 products on Datadog GovCloud, including Cloud SIEM and cloud security management.

Next, we landed a 7-figure annualized deal with a large American financial services company. This customer has a very seasonal business and experiences dozens of major incidents during the annual peak season with an average downtime per incident of about 5 hours. And they estimate millions of dollars of lost revenue for each hour of downtime. By replacing its cloud providers monitoring tool with Datadog and in particular, using our Real User Monitoring product, this customer targets substantial reductions in downtime.

They are starting with 5 Datadog products and are trialing our network monitoring, database monitoring, cloud security and cloud cost management products as they look to consolidate dozens of homegrown commercial tools.

Next, we landed a 7-figure annualized expansion with a major airline in Europe. This customer has adopted Datadog for its customer-facing websites. They are now moving hundreds of applications from on-prem to AWS, and they want to derisk their cloud migration. They estimate that each incident can cost tens of millions of dollars in lost revenue and customer impact. By using Datadog across 5 products, this customer expects to significantly improve mean-time-to-resolution, and they have already seen progress in that respect during their evaluation period with Datadog.

Next, we signed a 7-figure annualized expansion with a division of a hyperscaler delivering next-gen AI models. This customer is very technically capable and already has a homegrown observability solution, which requires time-consuming customization and manual configuration. They will be launching new features for their large language models soon and need a platform that can scale flexibly while supporting proactive incident detection. By expanding the use of Datadog, they expect to efficiently onboard new teams and environments and support the rapidly increasing adoption of the LLMs.

Next and last, we signed a 7-figure annualized expansion with a leading online food delivery company in Latin America. Before Datadog, this customer suffered from excessive alerting noise, siloed teams and lack of visibility with each minute of downtime resulting in thousands of lost orders. By using Datadog, this customer has experienced meaningful reductions in mean-time-to-resolution and false alerts while saving on hard costs in the Kubernetes environment. This customer is expanding to 10 products in the Datadog platform. And that is it for another productive quarter from our go-to-market team.

Now, let me say a few words on our longer-term outlook. Overall, we continue to see no change to the multiyear trend towards digital transformation and cloud migration, which we continue to believe are still in early days. We are seeing continued experimentation with new advances such as next-gen AI, and we believe this is one of the many factors that will drive greater use of the cloud and other modern technologies. So we are helping our customers every day to observe, secure and act on their business-critical applications and workloads.

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

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

Thanks, Olivier, and good morning. Q3 revenue was \$690 million, up 26% year-over-year and up 7% quarter-over-quarter. To dive into some of the drivers of our Q3 revenue growth. Overall, we saw trends for usage growth from existing customers that were consistent with our expectation. We've seen conditions remain roughly stable throughout 2024 with continued movement to cloud and modern DevOp technologies and with customers remaining cost conscious and seeking efficiency and value from their spend.

In Q3, we saw usage growth from existing customers that was higher than usage growth in the year ago quarter as well as higher than usage growth in the prior quarter. Now some of our growth is coming from AI native customers who this quarter represented more than 6% of our Q3 ARR, up from more than 4% in Q2 and about 2.5% of our ARR in the year ago quarter.

AI native customers contributed about 4 percentage points of year-over-year growth in Q3 versus about 2 percentage points in the year ago quarter. While we believe that adoption of AI will continue to benefit Datadog in the long term, we are mindful that some of the large customers in this cohort have ramped extremely rapidly and that these customers may optimize cloud and observability usage and increase their commitments to us over time with better terms. This may create volatility in our revenue growth in future quarters on the backdrop of long-term volume growth.

Now regarding usage growth by customer size in Q3, similar to the last quarter, we saw the strongest performance among our largest customers who spent multiple millions of dollars with us annually. And as we look at usage growth by segment, similar to recent quarters, we are seeing the strongest growth from our enterprise customers, where year-over-year growth in usage has accelerated over the past several quarters. Meanwhile, our SMB customers remain solid with year-over-year growth similar to the past several quarters.

As a reminder, we define enterprise customers as our clients with 5,000 employees or more, mid-market customers with 1,000 to 5,000 employees and SMB customers as those companies with less than 1,000 employees. Regarding our retention metrics, our net revenue retention percentage was in the mid-110s in Q3 with continued improvement from last quarter. This is a trailing 12-month measure. Meanwhile, we've continued to see an increase in recent quarters as we look at the NRR quarterly trend. And finally, our trailing 12-month gross revenue retention percentage remains stable in the mid- to high 90s.

Now moving on to our financial results. First, billings were \$689 million, up 14% year-over-year. Billings duration decreased slightly year-over-year. Pro forma for changes in billing timing and the slight change in duration, billings growth would have been in the mid-20s percent range. Our billings and billings growth can be volatile on a quarterly basis depending on the timing of our deals. Our 12-month trailing billing growth is similar to our trailing 12-month revenue growth with both in the mid-20s percent.

Remaining performance obligations, or RPO, was \$1.82 billion, up 26% year-over-year. And current RPO growth was in the high 20s percent year-over-year. RPO duration was down slightly year-over-year. Normalizing for duration, RPO growth would have been in the high 30s percent year-over-year. We continue to believe that revenue is a better indicator of our business trends than billing and RPO as those fluctuate on a quarterly basis relative to revenue based on the timing of invoicing and the duration of customer contracts.

Now let's review some 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 \$560 million, representing a gross margin of 81.1%. This compares to a gross margin of 82.1% last quarter and 82.3% in the year ago quarter. Our Q3 OpEx grew 21% year-over-year, the same growth as last quarter, though it would have been -- although it would have represented an acceleration from last quarter, excluding the impact of our DASH User Conference in Q2.

As we said before, we are investing in headcount in 2024, and the acceleration in OpEx reflects our execution on our hiring in sales and marketing and R&D so far this year. Q3 operating income was \$173 million or a 25% margin compared to 24% last quarter and 24% in the year ago quarter.

Turning to the balance sheet and cash flow statements. We ended the quarter with \$3.2 billion in cash, cash equivalents and marketable securities. Cash flow from operations was \$229 million in the quarter. And after taking into consideration capital expenditures and capitalized software, free cash flow was \$204 million for a free cash flow margin of 30%.

Now for our outlook for the fourth quarter and for fiscal 2024. First, our guidance philosophy remains unchanged. As a reminder, we base our guidance on trends observed in recent quarters and apply conservatism on these growth trends. For the fourth quarter, we expect revenue to be in the range of \$709 million to \$713 million, which represents a 20% to 21% year-over-year growth rate.

Non-GAAP operating income is expected to be in the range of \$163 million to \$167 million, which implies an operating margin of 23% and non-GAAP net income per share is expected to be \$0.42 to \$0.44 per share based on approximately 361 million weighted average diluted shares outstanding. And for the full fiscal year 2024, we expect revenue to be in the range of \$2.656 billion to \$2.660 billion, which represents 25% year-over-year growth.

Non-GAAP operating income is expected to be in the range of \$658 million to \$662 million, which implies an operating margin of 25% and non-GAAP net income per share is expected to be in the range of \$1.75 to \$1.77 per share based on approximately 359 million weighted average diluted shares.

Now finally, some additional notes on guidance. We expect net interest and other income for the fiscal year 2024 to be approximately \$140 million. Next, we expect cash taxes in 2024 to be in the \$20 million to \$25 million range, and we continue to apply a 21% non-GAAP tax rate for 2024 and going forward. Finally, we expect capital expenditures and capitalized software together to be in the 3% to 4% of revenue range for fiscal 2024.

Now to conclude, we are continuing to execute on our strategy, investing in our innovation and expanding our platform to deliver more value to our customers. And lastly, I want to thank all Datadogs worldwide for their efforts as we close out 2024. And with that, we'll open the call for questions. Operator, let's begin the Q&A.

# Question and Answer

## Operator

[Operator Instructions]

The first question comes from Mark Murphy with JPMorgan.

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

Congrats on another very healthy performance. Olivier, we noticed your AI contributions surged to about 6% this quarter. We're watching all the advances in the foundation models, including the OpenAI's strawberry version. We're watching the multistep reasoning, how they're becoming multimodal, the longer duration inference and video models.

Does it seem reasonable to you that we are on the cusp of a pretty interesting period in the next couple of years for the inferencing loads and that might drive some incremental traction for Datadog that is tied to AI? And then I have a quick follow-up for David.

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

I mean, look, it's definitely a very interesting period to say the least. We see tons of innovation across our customer base, still largely more around experimenting and testing new applications. So as we reported, we are seeing some customers moving to production, and we are seeing our production-minded LLM observability products, for example, being used by real paying customers with real volumes and real applications in real production workloads. So that's exciting and healthy. I think it's a great trend for the future.

In general, in terms of the workloads, you're right that we're starting to see more inference workloads, but they still tend to be more concentrated across a number of API-driven providers. So there are a few others, both on LLMs and other kinds of models. So this is where I think most of the usage in production at least is today. We expect that to diversify more over time as companies get further into production with their applications and they start to be customizing more on their models.

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

Okay. Understood. And then, David, to the extent -- and you've always said that revenue is a better indicator than billings and RPO. But to the extent the billings growth was affected by timing, as you mentioned, we've seen that before. We know it can bounce around. But did you have some invoices that would have gone out in September and instead, they were issued in October? In other words, would we see some recapture of the timing element in Q4 or perhaps in early next year? Or is it some other dynamic there?

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

Yes, it's really some other dynamic. It's that the timing of billing for last year was slightly different than for this year. And so it was more a factor of timing for billing last year that didn't repeat this year. So we think that the weighted average, what we talked about, the average over the 12 months is a better indicator of the relationship between billing and revenues. And when you look at that, they're much more closely aligned.

## Operator

The next question comes from Sanjit Singh with Morgan Stanley.

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

Olivier, in the framework that you sort of laid out for the business, Observe, Secure and Act, I wanted to focus on the last 2 pillars, Secure and Act. When we think about how the security -- cloud security sales motion has been going this quarter versus prior quarters. Any trend lines there? And any sort of early indications on the uptake of some of the service management products and more of the automation features within the Datadog platform?

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

Copyright © 2024 S&P Global Market Intelligence, a division of S&P Global Inc. All Rights reserved.



Yes. So on the security side, I think we -- there's quite a bit of focus right now on the cloud SIEM as we see a number of very exciting opportunities there. I think when you look at the landscape competitively what other companies are using, because most companies have a SIEM already at least for some part of their business, I think there's a very interesting opportunity for us there. And the product is mature enough to win in best-of-breed situations across the existing products. So we are making quite a bit of a push.

And you know we're still investing quite a bit in the rest of the platform and bringing all the products together to be a unified platform for security. But I would say the tip of the spear for this quarter and next is slightly -- I think there's a very specific opportunity there. On the -- and we also have a few other products security that are coming online soon that we haven't brought to GA yet, but that we think can make a big difference.

In the service management, we actually see very exciting trends from customers. We mentioned on the call, our OnCall product, which puts us directly to the paging loop and the start of many incidents. The product is getting stronger reception than we initially thought to the point where customers are basically clamoring to buy even though it's still in preview. So we feel good about handling the full loop for them, starting with when we detect something in observability, all the way to full incident resolution and automation as a big part of it.

So we think OnCall can be a bit of a watershed for us to do quite a bit more on that. So we're also excited, and we're doing what we can to accelerate the road map there because we think there's a very good opportunity. We have a few other building blocks on the service management that we released over the past year. And they're all growing quite nicely, I would say, from a small basis today, but we're very excited in the way they all come together to form an integrated platform where we can fully automate resolution for our customers.

So that's, again, still something we're building largely from those building blocks, but something that we think is going to be very exciting in the quarters and years to come.

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

That's great to hear. And sort of one follow-up going to sort of the spending intentions from your customers. It feels like most of the past 12 to 18 months, the sales playbook, particularly in the enterprise has been around consolidation. Is that still the theme in terms of driving new bookings, new expansion deals? Or are you starting to see customers focus more -- less about the consolidation opportunity and more about innovation and bringing their innovation budgets to bear to invest in AI and cloud and bringing on Datadog along for the ride?

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

There's always been innovation and new things. You're right that at least for the deals we talk about in the earnings, there tends to be more consolidation, but that's the nature of it. I mean innovation typically happens gradually as opposed to being a big bang customers switching \$5 million from 5 vendors into another vendor. So we talk about these less in those earnings calls, but that's been happening throughout.

We definitely see room for a lot more consolidation moving forward, both in terms of existing customers and new logos. And at the same time, we are excited to see what's happening with the AI innovation as it gets further down the pipe and away from testing and experimenting and more into production applications. And we have some signs that it's starting to happen. Again, we see that with our LLM observability product. We see that also with some of the workloads we monitor from our customers on the infrastructure side. But I would say it's still very early days in terms of customers being in production with their next-gen AI applications.

**Operator**

Our next question comes from Raimo Lenschow with Barclays.

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

Congrats from me as well. Can I stay on that subject a little bit longer, Olivier? The -- if you think about it, the one big driver for you guys in the past was workload growth, and that's something that we're kind of all watching out for the hyperscalers as well, but it looks like there are stable trends at the hyperscalers and for you as well.

Like if you think about the puts and takes there to think that, that's changing, like how much of that is just plain vanilla macro? And how much of that is like taking project resources away, not so much money, but also time for kind of AI, and it's obviously like, as you said, early in its life cycle. Can you just -- how do you see that playing out? Because that's the one thing we're all watching out.

**Olivier Pomel**

*Co-Founder, CEO & Director*

Yes. I mean, look, the key thing to remind everyone of is we think the general move of workload growth in the cloud or cloud environment, it might be public cloud-based, but they might also be private cloud. So this growth and move to the cloud is going to last for a very long time and it's going to be at a high -- fairly high level for a very long time.

So you're right that when we look at the numbers from the hyperscalers and when we try to factor out what's GPU driven, the growth looks stable-ish. We think that growth is still high and is going to last for a very long time, and that's one of the big underpinning trends that we're going to ride for the years to come.

Now on the edge of that, you're right that the -- where the workloads could have grown maybe instead of growing 20%, they could grow 25%, maybe some of those 5% instead are being invested both in terms of infrastructure budget or innovation -- time innovation budget. All that is going into AI, and that's largely right now in experimentation and model training and that sort of thing. That's probably right. But we see that also as a precursor to further workload growth in the future, like I would say, more traditional production application workload growth. So for us, it doesn't really change the equation.

It does create a bit of a decorrelation between our numbers as a whole and the numbers you might see reported from the hyperscalers, where a lot of their short-term growth rides more into -- on the capacity of GPUs they bring online for those experimentation and training workloads.

**Raimo Lenschow**

*Barclays Bank PLC, Research Division*

Okay. Perfect. And then one follow-up for David. Like if you think about capacity, like where are you in terms of sales capacity in case things are changing? Or how do you think about like increasing, maintaining capacity as we think about the new year?

**David M. Obstler**

*Chief Financial Officer*

We think long term, our revenue -- our sales capacity is highly correlated with our top line growth. And as we said, we've been attempting to increase our sales capacity similar to top line growth rates. And that has to do, as we've mentioned, with bottoms-up planning, putting sales capacity in areas where we see either under covered for a lot of white space and then ramping and training our salespeople. So our philosophy is to scale our sales capacity roughly in line with the top line.

**Olivier Pomel**

*Co-Founder, CEO & Director*

We see plenty of opportunity for growth and we're like -- the plan is to keep growing the sales capacity. There are many markets, many geographies where we are underpenetrated. We mentioned India on the call, it's one of them where there's tremendous opportunity, and we have very little presence at this point.

And in general, there's a question of what we're doing strategically, where we're going and how much we want to invest. And then tactically, the question of how we execute against that plan. And it's -- that one is a lot harder because it's easy to say we're going to grow the sales team by 20%, 30%, 40%. It's much harder to then have these people show up in the right region at the right time and be trained in everything else.

And I would say we're doing fairly well there. I would say we probably are doing better at the end of the year than we did at the beginning of the year in terms of bringing all of our -- growing our sales capacity for tactical reasons. But overall, we are executing towards our plan.

**Operator**

Our next question comes from Kash Rangan with Goldman Sachs.

**Kasthuri Gopalan Rangan**

*Goldman Sachs Group, Inc., Research Division*

I'm curious to get your thoughts on 2 things. I'll keep this brief. In the subsequent few weeks after the quarter closed, I'm curious, especially with rate cuts, how you feel the customer feedback has been coming along, particularly on the SMB side since the rate of change has not been discernible at least as of the end of the quarter.

And one for you, Olivier. As you look at GPU workloads, what about the company's existing portfolio geared towards CPU applies in a GPU world? And how do you monetize an instance of GPU versus a CPU if that were to make sense?

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

You go GPU?

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

I can go GPU, yes. So look, there's 2 parts to it. There is -- how do you -- what can you do for GPU that's different from a CPU. I think in general, there's quite a few things that are the same in terms of understanding how that machine sits with the rest of the application or like how the system performs and that sort of stuff. And there's quite a bit that's new and differentiated around profiling in general and understanding how you maximize the usage of the GPU bandwidth, which is usually what it's about.

I would say we're -- so right now, we're working with a number of customers that have real-world large inference workloads on how we can help on the GPU profiling side for inference. We're doing less on the training side, mostly because the training jobs tend to be more bespoke and temporary, and there's less of an application that's attached to those that these are just very large clusters of GPUs. So it's closer to HPC in a way than it is to traditional applications, though we are also experimenting with what we can do there.

There is a world where maybe in a durable fashion, 60% of workloads are inference and 40% are training. And if that's the case, there's going to be a lot of value to be had by having repeatable training and repeatable tooling for that. So we are also looking into that.

Today -- as of today, we really don't monetize GPU instances all that well compared to the other CPU instances. So GPU instance is many times the cost of a CPU instance, and we charge the same amount for it. That doesn't have to be the case in the future. If we do things that are particularly interesting and it's going to have a real impact on -- and deliver value and how customers use and make the best of their GPUs and in the end, save money. So today, that's the current situation.

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

And as to your other question, in the period since the closing of Q3, we've seen very similar trends to the year-to-date in Q3 in terms of customer growth, strength in enterprise, stability in SMB. And we have a strong pipeline for Q4 that we're working hard on harvesting as we get towards year-end.

**Operator**

Our next question comes from Brent Thill with Jefferies.

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

David, on RPO, it has been decelerating. It's in focus with investors. I know you've said focus on revenue. But I think many are kind of concerned about the trajectory. And I know you've said, hey, the backlog is strong and you're adding sales capacity and all the things that makes sense that your backlog is maybe bigger than the numbers you're reporting. But can you add any more context to what the pipeline -- how you feel about the pipeline? What you think is going on with the metric in terms of the decel? It was a pretty big step down.

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

Yes. I think -- as we started talking about duration, we had a surge in this period, the period last year of longer-term contracts. That was really customer-led. We may well have that in the future. But in just comparing the timing of that, it doesn't affect revenues. We had that kind of compare, I would say. When you look at adjusting for duration and looking at over a long term, we do find that all of these metrics are circling around revenue in the mid-20s.

So I think there's a lot of noise in those, and it has a lot to do with the timing, particularly in RPO of multiyear deals. We'll report at the end of Q4 how that progressed. We have a lot of business to do in Q4, particularly in the larger customer and enterprise side. And we'll see where that lands. But I think it's, again, best to go back towards revenues as the North Star.

**Olivier Pomel**

*Co-Founder, CEO & Director*

And those really don't relate all that much to the sales pipelines because a lot of those -- I'm putting my CFO hat right now, but -- which I never wear. So bear with me. Everything has to do with when customers recommit and extend their commit and for how long they extend this commit, and there's a high degree of variability there. So they might go from 1 year to 3 years. Those 3 years might actually run out or they might recommit after 1.5 years or 2.5 years or 3 years or maybe 3 years and 2 months if it takes longer to figure out what level they want to be at next.

And so it has a ton of variability to the billings numbers in general. And that's why we don't look at them at all for managing the business. We look at usage, and we look at the sales pipelines in terms of what new business units, what new workloads, what new revenue and usage are we getting from customers.

**David M. Obstler**

*Chief Financial Officer*

I think what Oli is saying is an output to the natural flow of business, which has more volatility than revenues still.

**Operator**

The next question comes from Karl Keirstead with UBS.

**Karl Emil Keirstead**

*UBS Investment Bank, Research Division*

Okay. Great. I've got a couple of questions about the comment, David, around some of your larger AI native customers, maybe revenues from them being a little bit more volatile going forward. I guess first question is why do you think that is? Is this just a natural phenomenon of a few of those AI native customers now becoming quite large. And so it's normal that they try to look for better unit pricing and optimization? Or is there something else going on?

And then maybe secondly, when you set your 4Q guide, did you embed the assumption of some of these AI start-up pricing and observability trends that you called out? Or no, you're just trying to be prudent and that might be a little bit more of a 2025 phenomenon?

**David M. Obstler**

*Chief Financial Officer*

Yes, do you want to?

**Olivier Pomel**

*Co-Founder, CEO & Director*

I can take some of that. I think the -- look, what we see there is we -- so we have that group of AI, like smaller -- relatively small number of AI companies or AI native companies. Many of them are model providers or infrastructure providers for AI that serve the rest of the industry and they are really a proxy for the future growth of the rest of the industry in AI.

That group has been growing very fast. I mean we mentioned it's 6% of our ARR. It was about 4% of our year-over-year growth in Q3 versus 2% one year ago. So it's been growing very fast. There is some revenue concentration within that group. So customers there follow a power law in terms of revenue like the rest of the customer base pretty much. And seeing this very fast growth, what we do expect, we've seen that before, is we expect optimization at some point and better terms at recommit from those customers because they're all way over -- or many of them are way over their last commit with us, which, again, goes back to the other comment on bookings, like their growth doesn't show up at all in the ARPUs and the bookings numbers because they're way over their commit.

The analogy I would give there is what we've seen with cloud native in the late '10s and early '20s, where we had these numbers of cloud-native consumer companies that were growing very fast, with 2 differences. The first one is that the AI cohort is growing faster and there are larger individual ACVs for these customers. And the second difference is that it's a much smaller fraction of our

total ARR. So at the time, we had a very large number -- a very large amount of our revenue that was cloud-native companies and consumer companies in the late '10s and early '20s. Today, we just have 6% of our ARR in that bucket.

So we made that comment because we don't have anything to say in terms of where it's going. We don't -- there's nothing we see in October that tells us like there's big changes. But we do think there can be volatility there. Especially, this can move the numbers in the short term, while the mid- to long-term growth -- or the mid- to long-term dominant motion we'll see is growth with that customer base. So that's why we made that comment.

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

Yes. And as to fourth quarter, I echo the comments we make overall, which is we take conservative assumptions relative to usage growth that are lower than we've seen in providing our guidance. Of course, where we are in the year, the effect on Q4 is more limited than it would be on next year. And as we report Q4 and give guidance, we will update everybody on any trends that we do see in this cohort at that time.

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

We didn't change guiding principles for this.

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

Yes.

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

We made an extra comment on it because we want to be transparent. We've seen that. And we see some customers are wherever they commit. And -- but we didn't make anything specific in the guidance.

**Operator**

Our next question comes from Kirk Materne with Evercore.

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

Oli, you mentioned the one federal customer this quarter -- larger federal customer this quarter. Can you just talk on broader trends for you all within the federal government and sort of the opportunity there and what you're seeing? And I guess, if it had any positive impact? Obviously, it's their fiscal year-end in September. So just any comments more generally on Fed would be great.

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

Look, it's a huge opportunity for us. And we're, I would say, quite early with some interesting successes. So we have some exciting logos there, like this one that we mentioned this quarter, which is an agency we all interact with and which has tremendous opportunity for growth and upsell with us.

There's another agency that already has a very large multimillion dollar ACV footprint with Datadog that we didn't talk about this quarter, but that's been a long-term customer. So there's plenty of opportunity. We're still very early in terms of capacity building, channel building and go-to-market in general in government.

So I would say it's still a small part of our overall business, but we see that there's tons of upside in it. And we're working hard on the product side so that we can capture that fully. So we -- over the past few years, we worked on our FedRAMP compliance, and we are working further on getting into more regulated, even tougher to get into workloads with FedRAMP and IL5 and other certifications like that. So we have a long road map there and big plans to build and a lot of upside.

**Operator**

The next question comes from Mike Cikos with Needham.

**Michael Joseph Cikos**

*Needham & Company, LLC, Research Division*

I just had 2 quick questions for you. The first on gross margin. Just wanted to get a better understanding, is there anything you can speak to as it relates to those AI native customers, the intensity of those workloads and then how your products feed into it or maybe even the broader portfolio and the product expansion you've seen. I'm just wondering if those newer products maybe -- are detracting from gross margin near term here versus some expansion that we can expect as these products scale? That's the first question.

Second question was just great to hear that the usage growth continues to trend higher, especially for those existing customers. Does it feel where we sit today in this "new environment" like this is par-for-the-course kind of growth when thinking about the usage coming from those existing customers? Or is there a reason to believe that this can actually accelerate if things can move higher one way or another? What would drive that?

**Olivier Pomel**

*Co-Founder, CEO & Director*

Yes. So on gross margins, in general, we are happy with where they are. I mean there's some small moves in it. I wouldn't read too much into moves. I mean there's -- behind the scenes, there's a number of things that underpin that. We keep building new functionality. At the same time, we keep optimizing our code and use it. We actually use our products to optimize that quite a bit. And then we also keep getting better and better agreements with our cloud providers as we scale.

So it's a combination of all that is what you see in the gross margin number. We -- internally, we have to make a call always on whether we direct more effort at building new functionality or optimizing. And the way we manage that is when gross margins get a little bit low, we put more effort on optimizing. And when we are in a happier zone, like we redirect more effort on new products and new functionality into that. There's not a lot of change or that product mix doesn't really matter all that much on a gross margin perspective. And so there's no particular worries there.

We think long term, there are plenty of opportunities to improve margins. But right now, the focus is on really shipping enough products to enough customers, making sure that those products get -- provide them with as much value as possible while staying within a certain happy zone on the overall margin.

In terms of the growth of workloads, look, I mean, as we said, we see growth across the customer base pretty much. We see growth of classical workloads in the cloud. We see large growth -- very large growth on the AI native side. We think that the one big catalyst for future acceleration will be those AI native applications or those AI applications, I should say, going into production for non-AI native companies for a much broader set of customers than the customers that are deploying these kind of applications to their -- in production.

And as they do, they will also look less like just large cluster of GPUs and more like traditional applications because the GPU needs a database, it needs core application in front of it, it needs layers to secure it and authorize it and all the other things. So it's going to look a lot more like a normal application with some additional more concentrated compute and GPUs.

**Operator**

The next question comes from the line of Gray Powell at BTIG.

**Gray Wilson Powell**

*BTIG, LLC, Research Division*

Maybe just a follow-up on Datadog on call. It was good to hear the commentary on that earlier in the call, and it has been coming up more in our field work. So I'm just curious like how should we think about the opportunity there? Is that something that could completely displace a product like PagerDuty? Or is it more of an add-on feature since my understanding is it mainly works with the Datadog ecosystem versus other tools?

**Olivier Pomel**

*Co-Founder, CEO & Director*

Look, we'll take it where the demand takes us. So we initially built it as a way for customers within our ecosystem to have a fully integrated experience. And really as a step stone towards full automation of incident resolution and full ownership from end-to-end. The part that strategically is the most interesting to us is the automation of the resolution, not physically aging customers.

That being said, the response from customers has been so strong that there is very high demand for integrating with many different other sources and plug ourselves into incident resolution loops that we might not have been a part of before even. And so that's definitely something that we're going to build for customers, and we're just very happy to see the demand there.

**Operator**

The next question comes from Ittai Kidron with Oppenheimer.

**Ittai Kidron**

*Oppenheimer & Co. Inc., Research Division*

Nice numbers, guys. Oli, a question for you. I think you mentioned in your prepared remarks that 15 of your 23 products are now running over \$10 million. Maybe if you look at the ones that are still under \$100 million -- or under \$50 million, where do you see -- where are you most excited? Which ones do you think have the highest odds of crossing the \$100 million mark?

**Olivier Pomel**

*Co-Founder, CEO & Director*

Well, again, I don't want to single out any products for which we didn't disclose metrics in particular. But look, we mentioned in previous calls, like there are products that are growing very fast that we think will reach the scale velocity. We see -- we talked about database monitoring, for example, was a product has been growing very fast, very clear value. And we -- I actually forgot which number we disclosed last time, but we did disclose a number for it. 1% of revenue is what we disclosed. And so this one is clearly headed for north of 50%.

There's a number in the other cohort that I -- a number of products across security, user experience that I think are definitely going to get there very soon. So we feel good about all that basically. The point here is pretty much every single one of those products should be above \$50 million. Some of them are going to get there faster than others. Some of them will cross \$100 million. Some of them will cross \$1 billion maybe. So I think we feel good about the product set.

**Ittai Kidron**

*Oppenheimer & Co. Inc., Research Division*

That's great. Maybe as a follow-up for both of you. You didn't provide '25 guidance, of course, but are there any thoughts you want to leave us with as we think about '25 and perhaps given where we are in the year. Oli, your initial conversations with customers and how they think about next year, anything to point out with respect to their behavior or investment areas of focus, which perhaps are different than what they've been in '24?

**Olivier Pomel**

*Co-Founder, CEO & Director*

Look, I think the only thing I would say is, I mean, I won't get into second-guessing our guidance or things like that. And we -- in general, also, it's very hard to guess usage ahead of time because that's actually very -- it is different or can be very different from the intents that are being manifested by customers or their understanding of what the next year is going to look like.

But the one thing I will say is we're investing. We're building sales capacity. We're definitely investing heavily in engineering. I think it's a target investment in the industry. Unlike many others, we don't expect at this point to have outsized investments in compute. We're not building absolutely large GPU clusters, but we are building engineering capacity, and we are building sales capacity. And so you should expect that in the numbers we give for next year.

David, anything to add or?

**David M. Obstler**

*Chief Financial Officer*

No, I think we -- I've said that at this point, we've noted that there's been stability to upward trend in usage that many of our clients, particularly in enterprise, are getting back to the work of launching digital applications, and that's creating the pipeline and the results for us, and although we still are in an environment that's careful and wants return on their investments.

And we'll update everybody as to whether that's what we see. But generally, we're giving comments based on what we see in this environment, which has been good and stable.

**Operator**

Copyright © 2024 S&P Global Market Intelligence, a division of S&P Global Inc. All Rights reserved.

[spglobal.com/marketintelligence](https://spglobal.com/marketintelligence)

This concludes the question-and-answer session. I would now like to turn it back over to CEO, Olivier Pomel, for closing remarks.

**Olivier Pomel**

*Co-Founder, CEO & Director*

Well, thank you all for attending the call. I want to give a few shout-outs. Firstly, for the product management team for building great products last quarter, the customers that actually spent all this time with us on getting those products to work right, especially for the new products, and I'm talking about LLM observability, for example, or OnCall, and so I know we spent a lot of time with customers to get those right.

And I also want to give a special shout out to the go-to-market team. We have a very, very, very loaded fourth quarter, very full slate for everyone in the next, let's call it, 1.5 months. And so I know everybody is super hard at work. So thank you, everyone. And on this, I think we'll wrap the call.

**Operator**

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



Copyright © 2024 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).  
© 2024 S&P Global Market Intelligence.