

TDC Earnings Call Transcript

Date: 2025-11-04

Operator: Good afternoon. My name is Charlie, and I'll be your conference operator today. At this time, I'd like to welcome everyone to the Teradata Third Quarter 2025 Earnings Call. [Operator Instructions] I would now like to hand the conference over to your host today, Chad Bennett, Senior Vice President of Investor Relations and Corporate Development. You may begin your conference.

Chad Bennett: Good afternoon, and welcome to Teradata's 2025 Third Quarter Earnings Call. Steve McMillan, Teradata's President and Chief Executive Officer, will lead our call today, followed by John Ederer, Teradata's Chief Financial Officer, who will discuss our financial results and outlook. Our discussion today includes forecasts and other information that are considered forward-looking statements. While these statements reflect our current outlook, they are subject to a number of risks and uncertainties that could cause actual results to differ materially. These risk factors are described in today's earnings release and in our SEC filings, including our most recent Form 10-K and in the Form 10-Q for the quarter ended September 30, 2025, that is expected to be filed with the SEC within the next few days. These forward-looking statements are made as of today, and we undertake no duty or obligation to update them. On today's call, we will be discussing certain non-GAAP financial measures, which exclude such items as stock-based compensation expense and other special items described in our earnings release. We will also discuss other non-GAAP items such as free cash flow, constant currency comparisons and 2025 revenue and ARR growth outlook in constant currency. Unless stated otherwise, all numbers and results discussed on today's call are on a non-GAAP basis. A reconciliation of non-GAAP to GAAP measures is included in our earnings release, which is accessible on the Investor Relations page of our website at investor.teradata.com. A replay of this conference call will be available later today on our website. And now I will turn the call over to Steve.

Stephen McMillan: Thanks, Chad, and thanks, everyone, for joining us today. Q3 marked another quarter of solid execution as we beat our revenue and recurring revenue guidance ranges. We delivered non-GAAP earnings per share of \$0.72, soundly ahead of our outlook, and we delivered free cash flow ahead of expectations. We posted our second consecutive quarter of total ARR growth ahead of our initial target of the fourth quarter. With a return to total ARR growth ahead of schedule, we have strong conviction in our durable growth path and expect this growth to continue in 2026. We also expect that the return to positive ARR growth, combined with the cost savings and productivity measures we've taken will result in meaningful free cash flow growth. Whether in cloud or on-prem, we are helping organizations build the data foundation and are delivering the enterprise context required for AI solutions. And we see the shift in our business from classic EDW towards the autonomous AI and knowledge platform. We see enterprises reevaluating how to cost effectively deploy Agentic AI. As we have noted for the past several quarters, we are seeing a resurgence of hybrid environments, which reflects a growing understanding of how enterprises can best leverage both on-prem and cloud capabilities. It isn't just about choosing between environments anymore. It's about effectively operating across both to meet diverse business needs. Our platform is designed to give customers the opportunity to run Agentic AI at scale, wherever that data resides in their business in public cloud, on-prem or private cloud. Interest in AI and in particular, Agentic AI continues to grow in virtually all industries. However, most companies are still in the early stages of deploying this technology and Teradata sits squarely at the center of this revolution. We believe we provide the enterprise context that AI agents need to deliver trusted, reliable results at scale. Without this knowledge, even the most advanced models can be just a plain wrong. This shift also creates a very specific opportunity because Agentic AI with its 24/7 always-on query potential can increase workloads on data platforms by up to 25x and use 50x to 100x the compute resources than what was required by previous modern analytic

workloads. Teradata is uniquely built to handle these mixed workloads and high volumes of tactical queries as enterprises deploy potentially thousands of agents and evaluate millions of relationships across thousands of tables to make a single decision, milliseconds matter. We not only manage the critical enterprise data that powers these AI systems, but we also can deliver the performance required at the level of performance and scale that AI needs. Teradata was built for these types of enormous workloads based on our massively parallel architecture, patented workload management and query optimization that is designed to provide a high-performance environment with predictable costs that can deliver the most complex AI workloads. Our patented QueryGrid data analytics fabric provides seamless high-performing data access, processing and movement across multiple data sources. Our industry data models are built on decades of working with the Global 1000. And through these, we bring deep context to language models, another area where we can bring unique benefits to our customers. We believe Teradata is the best autonomous AI and knowledge platform for Agentic workloads and that our platform provides the best price performance, whether on-prem or in the cloud. In the quarter, we were named a leader in the Forrester Wave Data Management for Analytics platforms, and the report noted that Teradata is a good choice for organizations seeking to support hybrid cloud DMA deployments, especially where reliability, scalability and high availability are essential. We're building the capabilities for the future to enable AI speed and scale. Earlier this year, we announced Enterprise Vector Store, a capability that enables organizations to include unstructured data and their integrated knowledge foundation. We also enhanced ClearScape Analytics with unified ModelOps capabilities designed specifically for Agentic AI. These provide seamless native support for open source models as well as CSP model APIs. We launched our MCP Server to deliver faster context autonomously, and we've recently taken several more significant steps to further our position. In September, we announced Teradata AgentBuilder, a suite of capabilities designed to accelerate the development and deployment of autonomous contextually intelligent AI agents. Now in private preview, it leverages open source frameworks, our MCP Server and deep semantic access to enterprise data across cloud and on-prem environments provided by our knowledge platform. Customers can develop their own agents or use ready-to-deploy Teradata agents to accelerate implementation and deliver rapid impact. Launched at our Possible event last month, Autonomous Customer Intelligence is a software and services offering that embeds Teradata agents across the customer experience or CX journey. These agents can uniquely leverage 4 decades of Teradata innovation and contextual knowledge from solving mission-critical industry-specific data challenges. Our integrated approach makes sure our agents are extensions of the enterprise data platform and broader knowledge ecosystem rather than generic tools that fail to deliver meaningful impact. To help customers transform AI pilots into production-ready Agentic solutions that deliver significant business value, we also launched new AI services. These new services are intended to make Agentic AI a reality at enterprise scale by combining embedded experts, proven methodology and Teradata's best-in-class autonomous AI and knowledge platform. Using a sprint-based use case-driven approach, Teradata AI services offer flexible tiered offerings that meet organizations at any stage of their AI journey from initial pilots to enterprise-wide Agentic deployments. Unlike competitors who offer either consulting or technology, we believe Teradata uniquely delivers both, enabling real-time context-aware agent decisioning that leverages our suite of AI tools, trusted data and decades of industry innovation. Working with our partners in an integrated approach accelerates deployment of autonomous intelligence, CX or otherwise to drive measurable business outcomes. We have forward deployed resources with deep expertise and talent. These AI/ML engineers and data scientists are working with customers across the globe, positioning Teradata as a leading AI/ML player and helping customers move from proof of concepts to production. This team is on track to complete more than 150 AI engagements with customers this year. We're also seeing a significant turn in our pipeline towards AI-fueled projects. Let's look at a few examples of wins from the quarter. These demonstrate the breadth of our offers in hybrid environments, cloud and on-prem. A multinational automotive manufacturer is expanding its Teradata Cloud platform on AWS to support increasing AI/ML workloads as it combat cybersecurity. Executing approximately 10 million SQL statements per day, the customer is moving beyond rule-based approaches and adopting AI/ML technologies to enhance its analytical capabilities. One of the largest U.S. health care providers deepened its strategic alliance with us as it further scaled its Teradata cloud deployment running on

Microsoft Azure. This expansion, building on momentum from earlier this year, underscores the provider's continued confidence in our high-performance cloud platform to support mission-critical data and analytics workloads. With this expansion, the organization is further positioned to drive operational excellence and harness complex health care data at scale across its entire system. A leading Japanese heavy industry manufacturer chose Teradata for its on-prem data platform as it transforms to a data-driven manufacturing entity and improves operational efficiency. A Central European financial services company recommitted to us through a 7-year partnership with Teradata as a Service on AWS. This enhances security, provides uninterrupted operations through disaster recovery systems that match production, supports monthly innovation testing and meet stringent data sovereignty requirements. We recently held our annual customer event named Possible. It was 3 days of high energy with our people, partners and customers speaking of what they are doing now with data and analytics and what they are looking ahead to do with AI and Agentic AI. It was our pleasure to recognize VodafoneThree, Ooredoo and Sicredi at the conference for demonstrating exceptional creativity, technical excellence and business impact through the use of AI on the Teradata AI and Knowledge platform. VodafoneThree in the U.K. was recognized for deploying an AI-supported fraud detection framework by leveraging AI to detect and mitigate fraud that has strengthened customer trust, improved regulatory compliance and enhanced operational resilience. Ooredoo Qatar, a leading Doha-based telco, earned this award for its advanced analytical capabilities and AI-powered customer engagement strategy. This strategy is built on Teradata VantageCloud and ClearScape Analytics, which were integrated with and run on GCP native services. Sicredi, Brazil's largest financial cooperative, was honored for its innovative use of ClearScape Analytics and our cloud platform to transform credit risk management as well as support sustainability initiatives. Most recently, Sicredi has also begun developing an AI agent to support provision analysis under Brazilian banking regulations, further strengthening its governance and risk management capabilities. We also held our first AgentBuilder workshop at the Possible event. This hands-on workshop was oversubscribed and packed with customers keen to build AI agents on Teradata. We're in the process of launching an online AgentBuilder experience to help accelerate the development and deployment of autonomous, contextually intelligent AI agents. It will be available from our website in the coming weeks. We held our annual partner forum concurrently with the Possible event, and we had strong year-on-year growth in partner participation. Companies that will win in the Agentic AI future will be the ones that create the most trusted interoperable foundation that lets every other AI innovation flourish. We believe that's our role in the ecosystem. We strive to be the trusted data foundation that makes everyone else's AI work better with the governance layer that lets companies experiment safely. We're partnering across all layers in the ecosystem, and we have strong partner co-sell activity in the third quarter, validating the strength in our ecosystem and identifying and nurturing new opportunities. While at our event, I hosted a fireside chat with one of our partners, ServiceNow. We discussed how together we can power autonomous operations at scale by combining our enterprise-grade analytics with ServiceNow's workflow engine. Our platforms work together to enable seamless integration, governance and automation. We're collaborating to help customers realize the full potential of their data, delivering intelligence and automation at enterprise scale. This is how we enable AI-native transformation for our customers, empowering organizations to break down silos, unlock real-time intelligence and transform every part of their business. By combining deep analytics, trusted data and intelligent workflow automation, we're enabling organizations to move from passive data collection to active Agentic operations, delivering real-time insights, proactive engagement and measurable business value. Exciting stuff, and that was just one of the leading partners that participated with us. We also hosted a number of industry analysts and a comment from Constellation Research summarize our focus on helping provide context to AI, noting that we believe there is no AI without context. That context isn't just data. It's the metadata, business logic and domain know-how that make AI decisions relevant and reliable. Without business context, even the best algorithms can't deliver the accuracy or explainability needed in real-world regulated environments. They also recognized that we are turning our decades of decision analysis experience into domain and industry knowledge models that give AI agents real context. And that our context intelligence framework captures how industries actually operate, so organizations don't have to start from scratch as we help teams build agents faster with

enterprise-grade performance, governance and trust already built in. Our hybrid capabilities are resonating in our customer base with interest in our recent product introductions, AI Factory, MCP Server and AgentBuilder, giving us further conviction that we offer a unique value proposition. We provide the flexibility to have consistent data, compute models, workloads, outcomes and experiences across a hybrid environment. We have full confidence in total ARR and are affirming our outlook for 2025. In our recent discussions with customers, we have seen how the Teradata Knowledge platform is ideally suited for AI workloads. AI is always on with ever-increasing agents driving massive complex query volumes. That's Teradata's sweet spot. Our ARR mix may vary as we see customers evaluating between cloud and on-prem for where to deploy the workloads as they build for their AI-enabled future. Regardless of the deployment options they choose, customers can rely on Teradata to run Agentic AI at scale and provide the context needed for trusted results. Thank you very much. Now I'll turn the call over to John.

John Ederer: Thank you, Steve, and good afternoon, everyone. I'm pleased with the progress we are making this year as we've demonstrated a return to consistent execution with our third quarter in a row of meeting or exceeding our guidance metrics. And perhaps as importantly, we expect that trend to continue in Q4 as we are reiterating our guidance for the full year. Looking at a few of the highlights for the third quarter. Total ARR growth was ahead of expectations, representing the second consecutive quarter of a return to positive growth. We exceeded the top end of our total revenue and recurring revenue guidance. We improved gross margin sequentially from Q2. We delivered considerable upside on our non-GAAP earnings per share, and we increased free cash flow on a year-over-year basis for Q3 and the year-to-date. Finally, as Steve commented, we are building a solid foundation this year to deliver continued financial improvement next year. In terms of our detailed financial results for the third quarter, total ARR grew 1% as reported and flat in constant currency. This is our second consecutive quarter of a return to total ARR growth, and this was driven by better retention and expansions in the quarter. At the beginning of the year, our target was to get back to positive total ARR growth by Q4, and we are pleased to be several quarters ahead of schedule. Cloud ARR grew 11% on an as-reported and constant currency basis, and the cloud net expansion rate was 109%. As discussed on our Q2 earnings call, we expected Q3 cloud ARR growth to be below our guidance range for the year due to the pull forward of a few deals last quarter. Total revenue was \$416 million, down 5% year-over-year as reported and 6% in constant currency, which was 1 point above the high end of our outlook due to higher recurring revenue. Recurring revenue was \$366 million, down 2% year-over-year as reported and 3% in constant currency, which was 1 point above the high end of our outlook. Recurring revenue as a percentage of total revenue was 88%, up from 85% in Q3 last year. Services revenue was \$47 million, which was consistent with our recent performance. We are seeing a transition in our services business this year as the team is moving from migration projects to delivering AI services, which we believe will provide improved performance next year. Looking at profitability and free cash flow. Please note that I will be referencing non-GAAP numbers for expenses and margins and a full reconciliation to GAAP results is provided in our press release. For the third quarter, total gross margin was 62.3%, which was up 70 basis points year-over-year. On a sequential basis, total gross margin was up 400 basis points, driven by improvements on both recurring and services gross margins. Recurring revenue gross margin was 68.9%, up 140 basis points sequentially. On services gross margin, we took actions last quarter to align our costs with current revenue, and we made substantial improvement in non-GAAP gross margin from negative 2% in Q2 to positive 8.5% in Q3. Operating margin for Q3 was 23.6%, which was up 110 basis points year-over-year and up 720 basis points sequentially. Overall, we are seeing improving margins as a result of cost efficiency actions we started last year. Non-GAAP diluted earnings per share were \$0.72, exceeding the top end of our outlook range by \$0.17. The outperformance was driven by higher recurring revenue and lower expenses -- we generated \$88 million of free cash flow in the quarter, which was up 28% on a year-over-year basis and provides us with increased confidence in our full year outlook. And finally, in the third quarter, we repurchased approximately \$30 million of our stock or 1.4 million shares. We continue to target returning 50% of our free cash flow to shareholders in the form of share repurchases this year. Turning to our outlook for the remainder of the year. For the fourth quarter of 2025, we expect recurring revenue to be in the range of minus 1% to minus 3% year-over-year on a constant currency basis. We expect total revenue to be in

the range of minus 2% to minus 4% year-over-year on a constant currency basis. And we expect non-GAAP diluted earnings per share to be in the range of \$0.53 to \$0.57. For fiscal '25, we reiterate our previous guidance for total ARR growth, and we are maintaining our range for cloud ARR growth. We have confidence in our total ARR target and continue to see a path to our cloud ARR range for the year. However, there are a handful of deals where customers are still assessing deployment options, which could have an impact on the mix between cloud and on-premise subscription ARR. We also reiterate our previous guidance for recurring revenue and total revenue. Given the guidance ranges that we provided for Q4, we anticipate recurring revenue and total revenue to be at the midpoint of our fiscal '25 ranges. On free cash flow, due to our strong performance year-to-date, we are narrowing the range to the top end of our initial outlook and now expect free cash flow to be in the range of \$260 million to \$280 million. Finally, we are raising our non-GAAP earnings per share guidance to a range of \$2.38 to \$2.42, reflecting our strong performance in Q3. Based on foreign exchange rates at the end of September, we anticipate 1 to 2 points of benefit to our Q4 '25 revenue. For the full year, we do not anticipate any material currency impact. Finally, we expect the non-GAAP tax rate to be approximately 23.1% and the weighted average shares outstanding to be 96.1 million for the full year. Again, please refer to our Q3 earnings presentation on our Investor Relations website for a complete list of our 2025 outlook ranges. In closing, we are taking actions that we believe will ultimately drive shareholder value. The first important steps were to: one, return total ARR growth to positive territory; two, focus on cost efficiencies; three, drive consistency in the business; and four, stabilize free cash flow, all of which are on track to achieve this year. As we start to focus on the objectives for next year, we are prioritizing our investments to capitalize on the substantial opportunity ahead for Teradata as a leading AI and knowledge platform for the autonomous enterprise. We believe that these investments, combined with the continued optimization of our business will enable us to deliver profitable growth and higher free cash flow. Thank you all for your time today. Now let's open up the call for questions.

Operator: [Operator Instructions] Your first question comes from the line of Erik Woodring of Morgan Stanley.

Erik Woodring: Steve, really nice to see the earnings and free cash flow upside this quarter. I think this was the first time since you began disclosing cloud ARR that we've seen sequential be negative intracalendar year for cloud ARR. I know you mentioned that it would dip below the target range this quarter, but I guess I look at the 11% and say it felt a bit below maybe where you would have expected sequential. But maybe you could just elaborate on how the quarter transpired for cloud ARR, when and where we see that net expansion rate bottom? And maybe why we just aren't derisking 4Q a bit, just given some of your commentary around customers assessing where they're going to be deploying with Teradata. And then a quick follow-up.

Stephen McMillan: Yes. Thanks, Erik. So I think on cloud, we did perform to our expectations. As we said in our Q2 earnings, we did expect that linearity to be below the full year outlook. I think as we look at the market, we're not seeing -- we're no longer seeing that kind of headlong rush to the cloud. It's much more of a nuanced decision of how our customers can accelerate the time to value for the AI workloads that they're deploying in their environment. And we're continuing to see that pattern of customers using our hybrid capabilities. I think I've said in the past that a good proportion of our cloud customers have deployments both on-prem and in the cloud with us. And they can decide across that massive data estate that they run from a Teradata perspective, where to run that workload. And I think that's nicely evidenced actually by the fact that our total ARR growth is ahead of schedule overall. We expect it to return to growth in the fourth quarter. So being ahead of schedule from that perspective and demonstrating that we're growing overall with our customers is a good achievement as we've executed through the year. We do start to see our net expansion rate starting to consolidate. And I think as we look at the overall results for the year in terms of our guidance, we were pretty confident in our total ARR growth. And as we look to 2026, we continue to see a path for continuing ARR growth in 2026.

Erik Woodring: Okay. And then just a quick follow-up. Your comment on -- I think you used the term meaningful free cash flow growth into 2026, at least to me, was the most confident I've heard you sound on free cash flow looking forward in a while. Can you maybe just unpack where this confidence comes from? I'm sure it has to do with ARR growth and some of your OpEx initiatives. But just -- I know you're not going to guide to 2026, but any way you can help us think about when you talk about

meaningful free cash flow growth, kind of what you're trying to tell us between the lines? And that's it for me.

Stephen McMillan: Yes. Thanks, Erik. I think you hit the nail on the head with the 2 points, but I'll just ask John to add any more color.

John Ederer: Yes. I think that's exactly right. I mean I think if you look at how this year is progressing, we've done a nice job on free cash flow relative to where we were at this point last year. And I think we're doing the right things and really focusing on this is a key driver for us. Certainly getting total ARR back to growth territory has had a positive impact. And then the cost actions that we've taken last year and this year as well are also supporting that number. And so we feel like we're putting the right pieces in place to continue that improvement next year.

Operator: Our next question comes from Radi Sultan of UBS.

Radi Sultan: First for Steve, last quarter, we talked about, I think it was roughly 1/3 of pipeline, including an AI component. So I guess, first, like how did that track this quarter? You mentioned the agent offerings, MCP Server. Like is there any area in particular within the AI portfolio moving the needle? And then any way to think about how these AI discussions more broadly are impacting competitive win rates?

Stephen McMillan: Yes. Thanks for the question, Radi. Yes, we're continuing to see that the AI influence pipeline increase. We saw it increase as we went through Q3, which is really great to see. We have supported that with a fantastic set of innovation and releases from a product perspective. Our new Chief Product Officer, Sumeet, is making a real difference there. In terms of we're measuring our innovation releases in terms of time from concept to press release. And I think as evidenced by the discussions we had with our customers at our most recent marketing event, they're really seeing those innovations is something that Teradata can provide in a holistic way to enable them to deploy Agentic AI workloads, whether it's from our Enterprise Vector store capabilities, our MCP Server, our AgentBuilder capabilities or ModelOps where we can include language model capabilities. So all of these, I think, are coming together. And one of our customers actually, I think, said it best where they said, Teradata is one provider in this area who's really putting it all together. But I think the most interesting thing, Radi, from a technology perspective is that we are seeing that the Teradata technology platform is really built for these AI workloads. When you think about an always-on AI agent, essentially that can execute thousands of queries and complex queries, so really large volumes of queries executing concurrently inside an environment with different types of workload. Our architecture inside Teradata, our massively parallel architecture, combined with our workload management and query optimization allows our customers to run those types of queries and the AI agents that they have developed to run those queries more effectively and efficiently than anybody else.

Radi Sultan: Awesome. And then second for John. You've been in the seat a couple of quarters now, strung together a couple of nice quarters. I guess is it fair to think about your approach to guidance being relatively consistent since you joined? And maybe are there any leading indicators or KPIs in particular that you're looking at that give you confidence in the outlook especially around Q4?

John Ederer: Yes. Thanks for the question. And I would say from an overall standpoint, I guess, in terms of guidance and our philosophy on it, we try to call it as we see it. And so we take a look at our forecast. We do have a number of KPIs that we'll look at from pipeline to our expenses to the revenue model, et cetera. There's a whole bunch of metrics that we'll take a look at and roll all of that up. When you talk about Q4, in particular, you're now getting down to the last few months, and we're literally going deal by deal. And so we've got that kind of granularity in terms of how we ultimately roll up the forecast and then our resulting guidance from it.

Operator: Our next question comes from Yitchuin Wong of Citi.

Yitchuin Wong: Steve, maybe I'll start with you. Great to see everyone in LA. I want to just follow on what Radi was asking, like around the competitive edge, like the Agentic AI strategy with autonomous at Possible and then AgentBuilder like really position Teradata like directly against some of this road map of your larger hyperscaler platform and even competitors like Databricks and Snowflake. Can you kind of help us understand what is really the longer-term durability competitive advantage that you see that Teradata can compete in the space? And is it more the hybrid cloud environment that you've been talking about and then the more enterprise IP within your decades of experience, maybe we can start

there.

Stephen McMillan: Yes. I think -- thanks for the question, YC. I think fundamentally, what sets us apart is actually our technology moat that we have already. It's the set of patented capabilities that allow us to execute these workloads in the most effective and efficient way, and we can do that both on-prem and in the cloud. So really being able to deliver that hybrid environment to our customers is clearly a differentiator for us. We announced earlier in the year our AI factory, which essentially combines a lot of capabilities together. We're partnered with NVIDIA in terms of the development of that AI factory. That gives us a fantastic base for future on-prem capabilities. We're looking forward to the next release of our Teradata technology platform, which will have GPUs built right into the platform in terms of executing that workload. But our customers are already using their Teradata on-prem platform to actually operate and execute AI workloads today in a very reliable way. And we're seeing it both in a real hybrid context, so both on-prem and in the cloud. At the end of the day, we see this as being a battle of the query engine. And we believe that our query engine is the best query engine to deliver AI-type workloads and to be that true knowledge platform in terms of building enterprise context for our customers. And that's built on all of the capabilities and solutions that we've developed over the last 40 years for our customers. We understand the domains around customer experience or supply chain management, and we understand the domains across all of the industries. So all of these things combined together to give us our unique positioning, and we are really excited about getting that message out to the marketplace and demonstrating that to customers on a day-to-day basis.

Yitchuin Wong: Well, understood. Maybe one for John here. You talked about like the much improved service gross margin to positive territory in the quarter. It's great to see an inflection there. Could you kind of help us double-click on some of these actions? I know we heard talk about the team starting to leverage more FTE or maybe even AI FTE here within the sales motion, especially with these newer AI use cases. Is this something that expected to go forward that could help drive better margin here and or even efficiency driving like AI, leveraging AI within the company that you touched on a little bit during [indiscernible] as well.

John Ederer: Yes. So I think a couple of different questions in there. I think in terms of the services business overall, -- to be perfectly honest, a lot of that is just rightsizing the organization for the current revenue stream that we've got there. We saw some headwinds this year due to higher migration activity in the prior year. And so for us, in the first couple of quarters, we were a little bit behind in that from a cost structure standpoint, and we fixed that in the second quarter, and we saw a nice rebound in Q3, and we think we've got some room to go in Q4. And so I think, again, there, it's just aligning the cost with the anticipated revenue. In terms of your broader question around margins overall and some of the things that we're doing with AI from an internal standpoint, there's actually quite a bit. And I won't do a justice, but I would encourage you to check out some of the presentations at our Possible conference. We had one that talked specifically about some of the things that we're doing internally. And there's a whole work stream around this that's really touching all parts of the business from cost of revenue through the operating lines.

Stephen McMillan: I'll just add to that, YC. I think from an AI services perspective, we're seeing customers have a real appetite to deploy real solutions. So with the launch of our customer intelligence framework and also backing that up with real consulting expertise, folks that can actually implement AI solutions inside our customers, we're really pivoting our consulting and services capability to deliver on something that we see as a supply-constrained marketplace in terms of folks that actually know how to deploy these solutions inside our customer base. So our most recent press release in the last couple of weeks around AI services and the capabilities that we have to help enable our customers in this market is super exciting. And obviously, of course, working alongside our partners to deliver those capabilities to the market is super important for us.

Yitchuin Wong: Yes, hopefully that traction continues.

Operator: Our next question comes from Chirag Ved of Evercore.

Chirag Ved: Following up on one of the prior questions here. I was wondering whether you could speak to the underlying trajectory of cloud versus on-prem at this point over the next couple of quarters, even qualitatively. Should we index more on the on-prem side of the business when we're looking out or -- and perhaps moderating cloud growth? And then any comments you might have on the associated

margin and pricing implications, if you can share that?

Stephen McMillan: Yes, I'll start and maybe John can make some comments. So I think -- Chirag, thanks for the question. I think we are definitely seeing that our on-prem is stabilizing, and we're seeing a rate of change and improvement. And that's due to and from an on-prem perspective, both retention and expansion of those on-prem environments. On that, we certainly are happy with our retention rates. It's in line with enterprise software and overall. But the fact is we'll take growth wherever we see it, right? And we're well positioned to take advantage of growth in this hybrid environment that some of our customers have got, but we're also really well placed to take advantage of on-prem growth for data sovereignty requirements or where the data gravity is on-prem. But we also see the fact that we can grow in the cloud successfully with our customers. We're seeing our expansion rates in the cloud pick up as we've gone through this year in comparison to some of our other years. And we expect that to continue from an expansion perspective. And so I think we've -- we're well placed to take advantage of the opportunity that's in front of us, and we'll see that growth in terms of that hybrid platform that we offer to the market.

Chirag Ved: Okay. That's really helpful. Maybe just one more. Great to see more of a focus on AI services within consulting. It really speaks to the importance and percolation of this technology. Looking ahead, do you see consulting revenue stabilizing a bit at this point, driven by the focus on delivering AI services? Or is this still a category that you're involved with, but starting to or continuing to shift over to your partner ecosystem?

Stephen McMillan: Yes. I think we've -- at our core, we're a technology company. We're about ARR growth. We do consulting and services to support our technology ARR growth. And clearly, the margin profile for that is where we want to operate. We've created that headroom, as we've discussed, in terms of creating that space for our partners to operate successfully with us. But I think every great technology organization needs a great consulting and services capability to support that technology value proposition. And it's great to be able to see our consulting and services organization pivot towards these AI services so that the relevance in the marketplace can increase and it can help our existing customer base and new customers that we come across deploy these AI solutions. And we see it actually as a great competitive differentiator. We've got a go-to-market motion now that supports a forward deployed engineering model to get POCs into our customers. But our consulting and services teams and their partners are going to ensure that they take those POCs from that proof of concept into reality and into production. And we've already seen success with our customers in terms of taking real problems and business domains that they have and turning it into production-ready capabilities. So really time to value from an AI perspective is super important, and we see our AI services capabilities is something that's going to support that.

Operator: Our next question comes from Matthew Hedberg of RBC Capital Markets.

Michael Steven Richards: This is Michael Richards on for Matt. Maybe just double-clicking on that dynamic where customers are assessing the deployment options. Just curious, is that a result of the announcement of the hardware refresh next year where maybe some customers are seeing the transformation you're bringing to the on-prem offering and now it's a bigger decision of whether or not to stay or move to the cloud? And then just any early feedback you've gotten on that decision to have this big refresh?

Stephen McMillan: Thanks, Mike, for the question. No, I wouldn't say it's got anything to do with the technology platform that we're coming out with next. I think the technology that we have in place today is actually enabling some of these decisions, both in terms of things like the AI factory, which are available today on the technology stack that we have. It's actually given our customers exactly what they want. They want the choice of deployment. They want to be able to choose where they put the workloads. And we offer our customers a workload first deployment model. So they can choose whether they want to run the workload in the cloud or whether they want to run it on-prem. And so that's the decision-making that our customers are going through. And the fact that we offer those technology capabilities in that hybrid environment is essentially given our customers the choice of deployment.

Operator: Our next question comes from Raimo Lenschow of Barclays.

Raimo Lenschow: Congrats from me as well on a great quarter. The quick question, Steve, more for you. If you think about the debate of where AI gets that data from, there is a kind of big debate kind of is

it coming out of the operational data stores and Oracle, et cetera, is making more out of like the data warehouses like you guys or more out of the data lakes. Can you speak to that, how you see that playing out? Or is it different use cases will have like a different data foundation?

Stephen McMillan: Raimo, you answered the question right at the very end. I think we are actually seeing customers want to get the best out of their data no matter where it sits. That's why we love QueryGrid as our technology to be able to combine all of these different data stores together. So no matter where the data is in the ecosystem, they can take that in a highly governed, reliable way and combine it together, whether it's coming from the data lake or whether it's coming from an enterprise data warehouse, and they can feed that into a language model in a very trusted environment. And that's what we're really delivering and offering to our customers. So I think this is all about -- if you think about AI solutions, they have to be trusted. They have to be ethical. You have to be able to track back through it, and they have to run efficiently and effectively. And that's what the Teradata platform enables our customers to do by combining all of those data sources together.

Operator: Our next question comes from Derrick Wood of TD Cowen.

James Wood: This is for you, John. You guys had nice outperformance on recurring revenue in Q3. But now for Q4, we had been kind of assuming low single-digit growth implied from your guide last quarter to now low single-digit decline. So was there any kind of pull forward of deals from Q4 into Q3? Or what would you call out on the change in the Q4 growth assumptions? And if I could just squeeze one other in on the cloud [ARR] having kind of dropped to 109%. Just remind us what the main drags to this number are? And any color on kind of when and where this could start to stabilize and perhaps move back up?

John Ederer: Sure. Yes. Thanks, Derrick. So on the recurring revenue side of things, I think our guidance for the year has actually been fairly consistent on the recurring revenue piece. We did have some variability if you look quarter-to-quarter, and that comes from the upfront portion of the on-premise subscriptions. And so depending on the mix of that in any given quarter, you might have more upfront revenue, which would otherwise throw off your expected linearity. In terms of the net expansion rate, we have seen some consolidation on that. If you look at what we've done historically, and even I think for this year, we're still on track for the same. About 50% of our expansion rate is coming from migration activity and the remainder is coming from expansions with existing customers. And so we see that continuing into Q4. Right now, you're seeing those rates consolidate. And so the net expansion rate is pretty close to what you're seeing for the cloud ARR growth overall.

Operator: Our penultimate question comes from Wamsi Mohan of Bank of America.

Wamsi Mohan: I think, Steve, you mentioned sort of cost takeout helping free cash flow into next year. Can you help us maybe think about just the absolute sort of OpEx trajectory going from here into '26? How are you thinking about the relative progression from here? And if I could, just -- I know federal is not really very large for you guys, but are you seeing any impact at all from the government shutdown?

Stephen McMillan: Yes, I'll take the first -- the last question first, Wamsi. No, we're not seeing any impact to our revenues as a result of the federal shutdown. And then just from an OpEx perspective, clearly, we've taken some fairly major restructuring activities through the year and a lot of them in the kind of June time frame and then into the September time frame. So we are expecting full year impacts and benefits to essentially ultimately, our free cash flow position as we move into 2026. So we are expecting that to amplify. And I think in relation to one of the other questions John hit it on the head, we're expecting that free cash flow growth to come from both our ARR growth expectations for '26 and also the operational efficiency, effectiveness, productivity measures that we've executed in 2025.

Operator: And our final question of today comes from Patrick Walravens of Citizens.

Patrick Walravens: John, this one is for you, too. And I know this was a good quarter, but if you divide your free cash flow by the revenue, you get like 21%. Not putting a time frame on it, I think where can that free cash flow margin go?

John Ederer: Yes. No, it's a good question, and I think it's somewhat related to what Steve just commented on in terms of the operating leverage. And if you look -- if you step back and look at what we did this year with revenue headwinds, we don't guide on operating margin specifically. But I think if you do the math and reverse engineer it, you're going to find that the operating margin has to be pretty flat and comparable with where we were last year. So that means in the face of revenue headwinds,

we're still able to capture that margin percentage. And we've done some things from an operational standpoint and a cost efficiency standpoint that will continue to benefit us next year. And so I won't give you a number today, but suffice to say that we've done some things this year that we think set us up well for next year from a margin and a cash flow standpoint.

Stephen McMillan: Thanks, Pat, for the question. And thanks, everyone, for joining us today. We are absolutely committed to show what the AI future holds for our customers and what our differentiated platform and capabilities can deliver. As we continue our focus on execution, we're really confident in our outlook, and we are looking forward to updating you all next quarter. Thank you very much. And operator, you can end the call.

Operator: Thank you. This concludes today's conference call. You may now all disconnect your lines.