Teradata Q3 2026 – Earnings Remarks

Chad Bennett, Senior Vice President, Investor Relations & Corporate Development, Teradata

Good afternoon, and welcome to Teradata's third quarter 2026 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. Please note that Teradata intends to file the Form 10-Q for the quarter ended September 30, 2026, 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, adjusted free cash flow and constant currency comparisons.

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.

teradata.

Steve McMillan Q3 2026 Earnings Prepared Remarks

Thanks, Chad, and thanks to everyone who's joining us today.

Teradata delivered a solid third quarter, performing as we expected across our operating metrics. We delivered another strong quarter of Adjusted Free Cash Flow growth, up year-over-year and year-to-date. And, we saw growth in Recurring Revenue, up 4% year-to-date, maintained strong operating leverage and exceeded our non-GAAP earnings per share range.

We are well positioned as we enter the fourth quarter. As in every year, the fourth quarter is our largest selling quarter, and we have already closed the largest expansion deal in our Q4 plan. We remain confident in our Total ARR guidance of 2 to 4%.

Today, I'll focus on the broad set of products and capabilities we are bringing forth as a result of the investments we made in our innovation pipeline. This was done with one goal in mind: helping customers leverage all their enterprise data and get their agentic AI initiatives into production. This is an imperative in order for AI to deliver durable value.

Organizations around the world face the challenge of accelerating AI and moving from experimentation into production. We believe that for AI to work autonomously, it needs to be trusted, ethical, and effective, and it needs enterprise context to be accurate. Organizations also have varied deployment requirements that are often shaped by regulatory and sovereign considerations. We are seeing increasing interest in our hybrid capabilities from customers running Teradata on prem, in the cloud, or in a combination of both.

As enterprises move from asking questions of the data to delegating work to AI agents, the platform underneath has to do more than it ever has before. Agents use data far more intensively, and very differently, than people do. We built the Autonomous Knowledge Platform to meet these challenges. And it is that backdrop for the announcements I will cover next, many of which are already available and being trialed by customers.

In September, we announced a major evolution of Tera, our agentic co-worker for enterprise data work. It is designed so that everyone from business analysts to database administrators can describe what they need in plain language, and see the work carried through without deep Teradata expertise. I personally use Tera in my daily work, so can attest that even business leaders can ask natural language questions and get the answers they need.

We built Tera to close two gaps that limit what enterprise AI delivers today. The first is business understanding. A model can be highly capable and still not know what a company's own terms, relationships, and policies mean. The second is execution, which has to be reliable, governed, and economical as agent activity grows.

Tera is generally available today, and we are extending it to take on both gaps, through Tera Context Engine, Tera Harness, and Agent Skills, all expected to arrive before the end of the year.

Tera Context Engine addresses the first gap by giving agents governed business meaning to work from. Our industry knowledge models are designed to add expertise validated with some of the world's most complex enterprises, so agents start with meaning rather than starting from scratch.

The second component is Tera Harness, which addresses the execution gap. It is designed to coordinate the right skills, tools, data, and models at each step of a workflow, with guardrails and human approvals built in to stop high risk actions before they run.

The third Tera element is Agent Skills, which take what Teradata has learned over decades and package it into functions that any agent can call… and that people can request in plain language. Supporting all three components is our AI Services team, which helps customers choose and implement the use cases where governed context will create measurable value.

We are very pleased with how Tera performed on SWE-bench Pro, a widely used software engineering benchmark. Using the same model for our internal testing, Tera used 73% fewer tokens than one of the leading AI coding agents, while achieving higher task completion rates and completing work 42% faster.

On data-engineering-bench, a public benchmark Snowflake Labs and Bespoke Labs developed for data pipeline engineering, Tera delivered 53% lower cost per reliably solved task than Snowflake Cortex Code. This result is based on our internal testing using the benchmark compared to Snowflake's own published results. Tera also earned the highest accuracy score on the benchmark in the same comparison. For customers, this level of efficiency reduces cost as a constraint on scaling AI, which gives them greater confidence to move more of their workloads into production on Teradata.

We're already seeing strong interest in Tera, with customers asking for proposals to utilize Tera across their entire environment for coding and data engineering.

We also released Tera internally to all employees so they can utilize its power, bringing a specific language model that they want to use in their own roles. Our internal usage helps our teams gain AI fluency and efficiencies, and provides near real-time feedback to our product organization on future enhancements to consider. Additionally, all of our engineering is now augmented by agentic development, and we've launched our first product built mostly using agentic development. We continue to expand how we use AI to accelerate how we deliver product capabilities while strengthening the quality and security of what we ship.

Two other recent announcements reflect our belief that AI should run where data already sits. In October, we announced the general availability of Elastic Python Compute in AI Studio. It gives data science teams a managed Python environment for scaling AI workloads in parallel, directly inside the Teradata platform. This offering is available on AWS and Azure today, as well as on-prem with Teradata Factory, so customers get the same capability in whichever environment holds their data.

September we announced an integration between our Teradata Autonomous Knowledge Platform and Microsoft OneLake. When running AI on Teradata, customers are able to use AI directly on data stored in OneLake without duplicating or migrating it.

For some customers, AI has to run entirely inside their own boundary, and interest in sovereign AI continues to grow. We think about sovereignty in three parts: control over where data lives and who can reach it, control over who operates the environment, and freedom to choose the models and technology an organization depends on.

At our most recent customer event in London, we expanded our sovereign offerings with two complementary deployment models.

The first is Teradata Factory with AI Studio, a unified environment that brings data, vectors, models, and agents into on-prem infrastructure. This is designed to give customers direct physical control over their AI execution.

The second is Teradata Sovereign Cloud on Google Distributed Cloud's air-gapped environment, for customers who want a cloud operating model inside an isolated boundary. It has no connection to the public internet, is operated by the customer, and is designed for governments, defense organizations, and regulated industries.

Control over where AI runs goes hand in hand with protecting the data itself. To this end, we also recently announced new privacy, security, and resiliency capabilities. They reflect how enterprise AI requires sensitive data to move between users, applications, models, and agents. This increases risk, and we believe protection has to work as a single system to safeguard the intelligence behind AI, without diminishing what makes it valuable.

With these new capabilities, customers can protect the identities behind their data while keeping the relationships that AI and analytics depend on. Post-quantum security and readiness strengthen protection against emerging quantum-era threats. And immutable backup protection keeps mission-critical operations available and recoverable. These are available now on Teradata Cloud across AWS, Azure, and Google Cloud, and masking and classification are also available on Teradata Factory and on-prem.

With the increased investment in our product pipeline, we've announced many new offerings that are already available, plus more are arriving this quarter and into early 2027. As agents take on more work, customers will need context, governance, and performance together, in whichever environment their data sits.

Technology alone doesn't get AI into production, so we pair our platform with AI Services. Our teams work alongside customers, starting from their data, their governance requirements, and their business priorities. We take use cases from prototype to production against clear business outcomes. We're seeing this across industries, from banking and aviation to telecommunications and pharmaceuticals. In each case, our teams are helping customers get AI into production and see value they can measure.

This year, we brought our new technology to the market at events around the world. These events encompassed product announcements, customer use cases, live workshops and demos - all showing how Teradata can help move customers' AI initiatives into production. Attendance beat our goals and our teams are now following up with customers. In addition, the workshops were oversubscribed and we are bringing the workshops online to extend the reach.

Now, I'll walk through some customer examples that reflect the breadth of our offerings.

- A multinational electronics company in Japan is a new logo for us. It selected Teradata for our hybrid capabilities, as the firm required both high-performance analytics across trillions of records and the stability of on-prem operations. Running on our on-prem platform and Teradata on AWS, this next-gen analytical environment will provide quality insights on its fast-growing manufacturing operations. Our team won by demonstrating superior query performance in a robust data foundation.

- An EMEA public-sector new account has selected Teradata to support a multimodal AI and analytics program for operational readiness and asset protection. This win reflects growing demand for trusted, governed AI platforms – platforms that can run mission-critical workloads on sovereign infrastructure, in highly secure on-prem environments.

- A federal transportation agency in Latin America selected Teradata on AWS, AI Studio and Elastic Compute to integrate fragmented traffic conditions, accident, and weather data. Teradata is also supporting AI-driven, predictive infrastructure decision-making. This win reinforces our role as a trusted data and AI foundation at scale.

- And finally, a utility company in the Americas chose our AI Services and AI Studio to help turn its trusted enterprise data into business value through AI. Building on its existing Teradata investment in governed intelligence and business context, it can reuse trusted data, metadata, business rules, and analytical assets - rather than creating a separate foundation for AI. This is to accelerate customer time-to-value with natural-language experiences while supporting scalable, governed AI adoption across the organization.

As we started the year, we said we would continue to drive top line growth, improve profitability and maintain healthy free cash flow. And, we'd invest to advance our products in helping customers derive value from AI. And we've done just that.

We are receiving positive feedback from customers on their AI initiatives with Teradata. Our more aggressive technology pipeline, coupled with enthusiastic response we've seen from customers, underpins our confidence in our annual outlook.

Now, I'll turn the call over to John.

teradata.

John Ederer Q3 2026 Earnings Prepared Remarks

Thank you, Steve, and good afternoon, everyone.

Q3 marked another quarter of progress against the goals we set at the start of the year. More specifically, entering the year, we were focused on driving Total ARR growth, continued operating margin improvement, and delivering strong free cash flow. On a year-to-date basis, we have had positive Total ARR growth every quarter this year, we have increased operating margin by 300 basis points, and we have grown adjusted free cash flow by 90%.

We also noted that we are focusing investments on product development this year. These investments are starting to pay off, and we are pleased with the accelerated pace of innovation this year. As Steve just discussed, we have advanced many elements of our Autonomous Knowledge Platform, and we are encouraged by the early response from customers. We believe these solutions position us well to capitalize on the market opportunity for Agentic AI as we head into next year.

In terms of our detailed financial results for the third quarter:

- Total ARR grew 1% as reported and 2% in constant currency, while Cloud ARR grew 11% as reported and in constant currency. As we mentioned throughout the year, we continue to anticipate the majority of our ARR expansion will come in Q4.

- Third quarter Total Revenue was $399 million and at the high-end of our outlook range, driven by recurring revenue.

- Third quarter Recurring Revenue was $359 million, which was also at the high-end of our outlook range. As we've discussed for several quarters, our revenue linearity this year has been skewed to the first half of the year due to upfront revenue recognition for on-premise subscriptions. So, while reported Q3 recurring revenue was down year-over-year, on a year-to-date basis, our recurring revenue is up 4% versus the comparable period last year. The timing of revenue this year has also had an impact on our quarterly margin performance, which I will highlight in a minute.

- Third quarter Consulting Services Revenue was $37 million, which was down year-over-year, as we transition this part of our business from cloud migration work to more advanced AI Services. During this transition, we have continued to optimize this part of the business to safeguard profitability.

Looking at profitability and 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 61%, which was down slightly versus Q3 last year. However, on a year-to-date basis, gross margin has improved by 140 basis points, due to a higher mix of recurring revenue.

- Recurring Revenue Gross Margin was 67%, down versus Q3 last year primarily due to revenue linearity this year.

- Consulting Services Gross Margin was 5%, up slightly from Q1 and Q2, as recent actions that we have taken are expected to have a bigger impact in Q4.

- Operating Margin was 22%, down versus Q3 last year, again, primarily due to revenue linearity. On a year-to-date basis, Operating Margin is 24%, up approximately 300 basis points versus the first nine months of 2025. Total operating expense has declined year-to-date, driven by lower SG&A, partially offset by higher investments in R&D to drive innovation.

- Non-GAAP diluted earnings per share were $0.70, exceeding the top-end of our outlook range by $0.11. The outperformance was primarily driven by lower operating expenses, along with revenue at the high-end of our outlook.

- We generated $96 million of Adjusted Free Cash Flow in the quarter, up 9% from $88 million a year ago. Year-to-date Adjusted Free Cash Flow is at $254 million versus $134 million for the first 3 quarters of 2025. As a reminder, Adjusted Free Cash Flow excludes all impacts related to the SAP settlement earlier this year.

- Finally, we continued to return value to shareholders, repurchasing $57 million, or approximately 2 million shares in the third quarter, bringing us to $131 million on a year-to-date basis, which is slightly ahead of our target of using 50% of our Adjusted Free Cash Flow for stock repurchases.

Now, turning to our annual outlook for 2026:

- We reaffirm our ranges for Total ARR and Adjusted Free Cash Flow.
- Given the year-to-date performance, we are narrowing our ranges for Total Revenue and Recurring Revenue to the high-end of our previous ranges. Total revenue is now expected to be in the range of minus 1% to 0% year-over-year. Recurring revenue is now anticipated to be in the range of 1% to 2% growth year-over-year.
- For non-GAAP earnings per share, we are increasing the range to $2.86 to $2.90.

For the fourth quarter of 2026:

- Recurring Revenue is expected to be in the range of -2% to -4% year-over-year.
- Total Revenue is expected to be in the range of -4% to -6% year-over-year; and

- Non-GAAP diluted earnings per share is expected to be in the range of $0.58 to $0.62.
- As we've discussed on prior calls, Recurring Revenue and Total Revenue growth rates in the fourth quarter are being impacted by the timing of upfront revenue recognition that occurred earlier this year. We believe our annual targets for Fiscal '26 Recurring Revenue and Total Revenue are more reflective of the general trends in the business.

In terms of some of the other modeling assumptions:

- For the fourth quarter, we expect the non-GAAP tax rate to be approximately 26% and the weighted average shares outstanding to be about 94.8 million.
- Also, we now anticipate FY26 other expenses to be approximately $17 million.

In summary, we are pleased with the financial and operational progress we have made over the first three quarters of this year, and we remain confident in our ability to achieve our full-year objectives. We also remain very focused on driving operating margin improvement. We now see operating margin expansion of approximately 200 basis points for the year which is double the 100 basis point improvement that we had anticipated at the beginning of the year. The Operating Margin expansion has driven our increased outlooks for both Non-GAAP EPS and Adjusted Free Cash Flow this year, with Non-GAAP EPS increasing from a midpoint of $2.60 at the start of the year to the current midpoint of $2.88, and the midpoint of our Adjusted Free Cash Flow range from $320 million at the start of the year to our current $340 million.

Thank you all very much for your time today, now let's open the call up for questions.