

Alibaba Group

September Quarter 2026 Results

Conference Call

Transcript

Thursday, 19 November 2026

Introduction

Lydia Liu

Head of Investor Relations, Alibaba Group

(Original)

Good day everyone, and welcome to Alibaba Group's September Quarter 2026 earnings conference call.

On the call with me are Joe Tsai, Chairman; Eddie Wu, Chief Executive Officer; Toby Xu, Chief Financial Officer; Jiang Fan, Chief Executive Officer of Alibaba E-commerce Business Group.

I would like to remind you that this call is also being webcast on our corporate website. A replay of the call will be available on our website later today.

Now I will quickly cover the safe harbor. Today's discussions may contain forward-looking statements, particularly statements about our business prospects and expected financial results, that are subject to risks and uncertainties, which could cause actual results to differ materially from those contained in the forward-looking statements. Please refer to the safe harbor statements that appear in our press release and investor presentation provided today. Please note that certain financial measures that we use on this call are expressed on a non-GAAP basis. Our GAAP results and reconciliations of GAAP to non-GAAP measures can be found in our earnings press release and investor presentation provided today.

Unless otherwise stated, growth rate of all stated metrics mentioned during this call refers to year-over-year growth versus the same quarter last year.

With that, I would like to turn the call over to Eddie.

Group Business Update

Eddie Wu

CEO, Alibaba Group

(Translation)

Greetings, and welcome to Alibaba Group's earnings call for the second quarter of fiscal year 2027.

This quarter, Alibaba Group revenue growth accelerated, and profitability continued to improve. Group total revenue grew 10% year over year. In AI Cloud and Compute Services, external revenue growth accelerated to 64%. At the same time, the year-over-year growth rate of recurring revenue from external customers increased to 67%. Adjusted EBITA margin rose to 14.5%. AI-related product revenue grew at a triple-digit rate for the thirteenth consecutive quarter, and its annual revenue run rate, or ARR, exceeded USD10 billion. Profitability improved significantly across our e-commerce businesses. In quick commerce, we continued to narrow losses substantially while maintaining market share.

The rapid improvement of AI model capability is accelerating the adoption of AI applications across industries and driving demand for training and inference. The deployment of AI agents at scale is further driving demand for compute. The combined result of these factors is that global compute demand will outstrip supply over the next several years. This profound demand is providing a solid foundation for AI commercialization at scale.

Next, I'll discuss four areas in more detail: AI and cloud commercialization, T-Head AI chips, models and AI applications, and our consumer business.

First, both revenue and profit growth in AI + Cloud are accelerating.

This quarter, revenue from external customers in AI Cloud and Compute services grew 64% year over year. That was the tenth consecutive quarter of accelerating growth. Adjusted EBITA increased 121% year over year, and adjusted EBITA margin expanded to 14.5%. This performance was driven by growth across compute, storage, MaaS, and AI applications, and the synergies among them.

AI-related products generated RMB17.9 billion in quarterly revenue and accounted for 49% of Alibaba Cloud's external recurring revenue. AI is now a core revenue pillar of our cloud business. The annual revenue run rate for AI-related products exceeded USD10 billion.

The strong revenue growth reflects broad demand for our AI + Cloud services. Beyond demand from AI labs for training and inference, we continue to see strong demand from customers

across various industries in China, including large state-owned enterprises as well as companies in intelligent manufacturing, automotive, embodied AI, and biotechnology.

Second, we continue to scale up T-Head AI chip volumes and strengthen our supply chain capabilities.

Starting with chips, T-Head has established a portfolio of proprietary chips, spanning the full stack of GPU, CPU, storage, and networking chips. Since launching in August, supernode instances powered by Zhenwu M890 have ramped up in volume and are now being commercialized at scale. We expect to begin mass production of the next-generation Zhenwu V900 chip in the first quarter of calendar 2027. The Zhenwu V900 delivers three times the computing performance of the M890. Next-generation supernodes powered by the Zhenwu V900 support large-scale network interconnections, enabling workloads such as training and inference for ultra-large-scale models. These supernodes will be the cornerstone of Alibaba Cloud's AI infrastructure. Zhenwu V900 supernode clusters will launch at commercial scale in the second quarter of 2027.

T-Head is accelerating its three-generation GPU roadmap. We fully rolled out Zhenwu M890 this year. We plan to ramp up the supply chain and move into large-scale production for the V900 in 2027 and the J900 in 2028. Infrastructure powered by our own domestically developed chips is becoming the foundation of Alibaba's full-stack AI capabilities. It will increasingly support model training, inference, and agentic workloads for both Alibaba Group and Alibaba Cloud customers, supporting growth in the group's AI revenue. Qwen is already training its next-generation model on T-Head's domestic chips, and Model Studio is already providing model inference on domestic chips. Our full-stack AI infrastructure will give us a clear advantage in the domestic compute supply chain. Deep integration between our proprietary chips and models will continue to lower the cost per unit of compute and help us gain greater efficiency in AI application development.

Taking all of these factors into account, our targets for calendar year 2030 remain unchanged, achieving USD100 billion in external cloud revenue with a margin exceeding 20%.

Third, we continue to advance our models and AI agent applications.

Our newly released foundation model Qwen4-Max has 2.4 trillion backbone parameters. We continue to advance its key capabilities, including agentic tasks and coding, where it's among the leading models globally. Top-tier partners have also recognized Qwen's capabilities. Apple Intelligence, which is deeply integrated with Qwen, is set to launch for users in China in December.

On the enterprise side, QwenWork launched in August and quickly became a leading enterprise AI agent. We recently launched QwenWork E-commerce, and the percentage of paying e-commerce users continues to rise. On the consumer side, the Qwen app is becoming a personalized assistant for getting things done. It now offers services for more than 20 daily-life use cases.

Fourth, e-commerce business's profitability grew significantly, with China E-commerce business contributed to stable operating profit.

This quarter, our e-commerce businesses delivered steady revenue growth and significantly higher profits. China E-commerce demonstrating stable operating profit. Our quick commerce business maintained its market share and continued to narrow losses, with unit economics improving quarter-over-quarter.

Two quarters ago, we began to see our AI investments cross the commercialization inflection point. Last quarter, growth accelerated, margins expanded. This quarter, those investments are delivering commercial results faster, and the returns are clearer. Our multi-year investment in compute and cloud infrastructure has given us full-stack product advantages and economies of scale which are becoming a key competitive advantage for Alibaba in the AI industry. They allow us to respond flexibly to industry shifts and offer AI products at scale to meet our customers' rising demands for AI. Looking ahead, we will continue to invest in AI infrastructure to extend our lead and strengthen our long-term competitiveness.

I will now hand over to Toby for the financial update. Thank you.

Financial Highlights

Toby Xu

CFO, Alibaba Group

(Original)

Thank you, Eddie.

This quarter, our business delivered strong growth in both revenue and earnings because of our strategic focus on AI + Cloud growth and profit sustainability in consumption business.

Revenue from AI Cloud and Compute Services grew 53% year-over-year, with adjusted EBITA margin sequentially rising to 14.5%. The growth of recurring revenue from external customers accelerated to 67% year-over-year. The strong growth reflected broad-based demand for our AI + Cloud services.

On consumption, adjusted EBITA of our E-commerce Group achieved 243% year-over-year growth this quarter, primarily driven by a significant loss reduction in quick commerce. In addition, our China E-commerce business contributed to stable operating profit and the Global Wholesale business delivered strong profits, while we reduced losses in the International E-commerce business.

Now let's look at the financial results for this quarter:

Total revenue increased 10% year-over-year to RMB271.5 billion, driven by the strong momentum in cloud business and quick commerce.

Total adjusted EBITA increased 260% to RMB32.6 billion, primarily attributable to the improved operating results in Taobao Instant Commerce and cloud business, as well as enhanced operating efficiencies across various businesses.

GAAP net income was RMB64.4 billion, an increase of 213%, primarily driven by the increase in net investment income and income from operations.

Operating cash flow was strong this quarter, increasing 55% year over year to RMB15.7 billion.

CapEx was RMB50.5 billion (USD7.5 billion) this quarter, due to continued investments in AI infrastructure to meet strong and accelerating customer demand. The significant year-over-year increase was due to an acceleration in scaling compute capacity, higher pricing of a broad range of chip components and increase in CPU-compute capacity driven by anticipated growing customer adoption of AI agents. The quarter-over-quarter decline from RMB67.7 billion (USD10 billion) in the previous quarter reflects fluctuations in procurement cycles. We will maintain our

disciplined approach by consistently tracking and evaluating the projected ROIC for future CapEx.

Free cash flow was an outflow of RMB34.5 billion, compared to an outflow of RMB21.8 billion in the same quarter last year. The decrease was mainly attributed to the investment in cloud infrastructure, partly offset by reduced losses in Taobao Instant Commerce, and increased customer prepayments for our cloud services. The quarter-over-quarter improvement from an outflow of RMB44.7 billion last quarter reflects our prudent cash management.

As of September 30, 2026, we continued to maintain a strong net cash position of USD32.9 billion, up from USD30.7 billion in the previous quarter. Excluding debt with maturities beyond five years, our net cash position was at approximately USD49 billion.

Our robust operating cash flow and net cash position, together with realized investment gains, give us the confidence and financial flexibility to invest in long-term growth.

We will maintain disciplined investments in AI capabilities with a focus on return on invested capital while targeting to achieve positive free cash flow.

Let me now turn to our segment results.

Alibaba E-commerce Group, we sustained revenue stability with CMR decreasing 5% year-over-year and growing 1% on a like-for-like basis if the contra revenue impact from the new business development program were excluded.

Revenue from China Quick Commerce business was RMB58.8 billion, an increase of 22%, primarily driven by Freshippo's robust growth across online and offline channels.

Alibaba E-commerce Group's adjusted EBITA grew 243% to RMB38.9 billion, primarily driven by a significant loss reduction in the China Quick Commerce business. Taobao Instant Commerce continued to improve its UE, reflecting higher average order value.

Marketplace e-commerce delivered strong earnings, supported by stable operating profit of the China E-commerce business, a loss reduction in the International E-commerce business, and strong growth in the Global Wholesale business.

Now let's review the business updates and results of AI Cloud and Compute Services.

Total revenue was RMB61 billion, a 53% increase from the same quarter last year.

Revenue from external customers was RMB47 billion, with growth accelerating to 64%. Recurring revenue from external customers, which excludes compute equipment solutions revenue, grew 67% to RMB37 billion. Customers continue to expand their use of our AI + Cloud services, driven by strong demand for AI computing and core CPU services, including storage, compute, security and networking.

This quarter's AI-related product revenue was RMB17.9 billion, implying an annual revenue run rate of RMB71.7 billion. It delivered the thirteenth consecutive quarter of triple-digit growth and accounted for 49% of external recurring revenue.

The adjusted EBITA margin expanded to 14.9%[1], benefiting from pricing enhancements amid tight market supply and scale-driven operating leverage.

We reiterate our calendar year 2030 guidance of achieving USD100 billion in external cloud revenue with a long-term margin exceeding 20%, driven by broader enterprise adoption and deeper integration of our AI products into core business workflows.

Next, let's look at AI Labs and Applications.

Total revenue was RMB4.4 billion, a 51% increase from the same quarter last year, primarily driven by rising demand for agentic AI products.

The segment narrowed loss quarter-over-quarter to RMB10.5 billion. We observed improved efficiency in both model training, and marketing spend on the Qwen app.

We are also continuously upgrading our AI agent offerings for both businesses and consumers through innovations across QwenWork and the Qwen app. In October, we introduced QwenWork E-Commerce, which is integrated into QwenWork and is able to automate end-to-end cross-platform operations for e-commerce merchants. Additionally, we have launched Qwen Go!, a cloud-based personal agent embedded within the Qwen app.

All others segment revenue decreased 5% year-over-year to RMB27.7 billion, reflecting the decrease in revenue from Cainiao.

All Others Adjusted EBITA was a loss of RMB3.3 billion, primarily due to our increased investment in technology businesses.

In conclusion, these strong financial results demonstrate that advances across our co-designed and integrated AI stack, together with the profit sustainability of our consumption business, position us well to capture increasingly diverse demand across industries and create greater value for our shareholders.

Building on this momentum, we will continue to balance disciplined investments for growth, with a focus on return on invested capital and proactive cash management to accelerate our journey to positive free cash flow.

Thank you, that's the end of our prepared remarks. We can open up for Q&A.

[1] Misspoken, "14.9%" should be "14.5%".

Q&A

Lydia Liu:

(Original)

Thank you, Toby. We will now begin the Q&A session. You're welcome to ask questions in Chinese or English. A third-party translator will provide consecutive interpretation. In the case of any discrepancy, our management statements in the original language will prevail.

Gary Yu (Morgan Stanley):

(Original)

My question is about the cloud revenue outlook. First of all, I appreciate the guidance of USD100 billion in 2030. But in light of the latest BIS announcement, how should we look at the cloud revenue growth in the next two quarters within the fiscal year?

And are we still confident to deliver over 50% revenue growth going into next fiscal year in FY 2028? If so, what makes us confident about it in light of the latest regulatory environment?

Eddie Wu:

(Translation)

In our overall planning for the 2030 long-term targets, our AI infrastructure has always been based on T-Head's domestic chip supply chain plan. As production capacity ramps up, domestic compute will account for an increasing share of our overall compute and will become the main engine supporting cloud revenue growth. AI compute remains supply-constrained globally, especially in China, where we expect supply to be particularly tight over the next couple of years. As we continue to advance and scale up our full-stack AI products and supply chain centered on T-Head, we continue to have strong visibility toward our previously communicated calendar year 2030 targets: USD100 billion in external cloud revenue and a profit margin above 20%.

As for next year, the new global export control rules on compute will in fact affect next year's products. I think there is very significant uncertainty across the entire global supply chain. However, our progress on domestic compute is ahead of our original expectations. So taking everything into account, there may be some impact on our operating plan and targets for next year. We are of course still evaluating this, but factoring everything in, we expect our operating target for the next fiscal year to fall within a 30% to 45% growth range. As our domestic

compute products continue to scale up in calendar year 2028, we expect to return to rapid growth of above 50% in that calendar year.

Kenneth Fong (UBS):

(Original)

I actually have a follow-up also regarding the tighter U.S. restriction affecting the access to advanced GPU and overseas compute. How should we think about the CapEx trajectory to reach our compute target of 20 GW? I understand the situation is still very dynamic. With a potential change in our CapEx cadence, would it be possible that we would turn free cash flow positive earlier than expected?

Eddie Wu:

(Translation)

I think compute policies across different regions of the world are indeed undergoing very significant changes right now, which is bound to have some impact on Alibaba Cloud's long-term data center build-out, both in China and globally. Based on our current preliminary assessment, the long-term 20 GW build-out target may be lowered, and our overall CapEx plan should also come down to some extent. However, we will have a more aggressive investment plan for the domestic compute supply chain.

Our preliminary assessment is that, as the compute supply chain centered on T-Head continues to scale up in calendar years 2027 and 2028, we expect to be able to turn cash flow positive in the second half of 2028. In other words, our heaviest sustained investment cycle will be in 2027 and the first half of 2028. However, after 2028, as we fully achieve large-scale, sustained volume expansion in the domestic compute supply chain, I think we will have passed through the heaviest phase of our AI infrastructure investment. So we remain very confident in our long-term cash flow planning.

Joyce Ju (Bank of America):

(Original)

My question is on your core e-commerce business. Given the soft macro environment and the cautious peer guidance for the fourth quarter, what is your December quarter outlook for Taobao Tmall in terms of CMR growth and adjusted EBITA trajectory?

On instant commerce, could you update us on your progress in the December quarter and your visibility on the timeline for further loss reduction next year?

Jiang Fan:

(Translation)

In the September quarter, the overall performance of our e-commerce segment was in line with expectations. Domestic e-commerce remained stable. Quick commerce maintained very good year-over-year growth while substantially narrowing losses. Profitability across our other businesses also improved meaningfully.

We have also observed that through the September quarter and into the December quarter, the domestic consumption environment has remained under pressure. Consumers have become more cautious in their spending, and certain categories continue to face intense competition. In the December quarter to date, our platform's key operating metrics have been consistent with the overall trend in China's online retail sales of physical goods. In the current consumption environment, consumer staples are showing better growth than discretionary products, and we will further strengthen our competitiveness in staples during regular selling periods.

Quick commerce will remain a core driver of platform growth over the next few years. In the September quarter, GTV of online non-food categories in quick commerce grew 40% year-over-year, leading the industry. In particular, our front warehouses expanded rapidly. In categories such as food, fresh produce, and daily necessities, quick commerce continues to reshape consumer habits and drive growth in these categories. On the food delivery side, our unit economics (UE) improved in line with expectations, market share remained stable, and average order value continued to rise.

In the December quarter, we expect quick commerce losses to continue narrowing. We continue to maintain our target of halving losses next year and expect to achieve profitability in quick commerce in FY29.

Beyond B2C e-commerce and quick commerce, other businesses within our e-commerce segment, including B2B and Xianyu, have also continued to make progress in recent years. This quarter, we integrated our international and domestic B2B businesses to better serve overseas demand with domestic supply. We believe this integration, supported by AI technology, will unlock even greater growth potential in the B2B business. In the September quarter, Xianyu achieved 29% year-over-year order growth and has now been profitable for consecutive quarters. We believe these businesses will deliver sustained revenue growth and profit contribution to the overall e-commerce segment going forward.

Robin Zhu (Bernstein):

(Original)

I guess there are two slight parts to my question. One is, could management discuss the impact of the BIS restrictions on model training and therefore the speed of iteration of the flagship Qwen models?

Related is, I guess our understanding is that a proportion of your data center capacity expansion is going to be funded by partners, T-Head customers. Could you share any updates on progress towards deals of this nature potentially as a proof of concept?

Eddie Wu:

(Translation)

The new regulations on access to compute are currently in the process of being implemented, and many detailed rules have yet to be issued. In our use of compute around the world, we have consistently complied with applicable laws and regulations. Therefore, we still need more specific information before we can make a relatively complete assessment of the specific impact of the new rules on our products and financials. That said, if the progress of some upcoming products is affected, I think it will be a relatively short-term impact, because by Q1 2027, our domestically produced T-Head chips will be able to fully support the training of ultra-large models.

In fact, we have already begun training our next-generation models on this year's generation of T-Head chips. As production capacity for T-Head's domestic chips continues to scale up, this will provide long-term support for our model development, because in the future the vast majority of models—including our own, as well as those of our partners and customers—will complete their entire training on domestic chips, including both pre-training and post-training. Therefore, as capacity continues to expand in 2027 and 2028, I think the pace of our long-term, continuous model iteration will actually become more controllable.

Ellie Jiang (Macquarie):

(Original)

I have a question mostly regarding T-Head. We are pretty encouraged to hear management commenting on the solid ramp-up pace for T-Head and obviously the promising pipeline. I am just wondering what proportion of the incremental training and inference capacity can T-Head supply over the next few years?

Also just a quick follow-up, can management comment on the latest updates regarding T-Head's spin-off and IPO timeline? Once this sort of comes out, how do we really assess the internal versus external resource allocation for T-Head, and what will be the implication on cloud margins?

Eddie Wu:

(Translation)

Based on the current situation, we expect that in 2027 and beyond, the vast majority of incremental AI chip supply will come from our domestic T-Head chips. We may also procure some other domestic chips to support our internal inference needs. Our own flagship chips will be prioritized for large-model training, first ensuring our own model iteration while also taking into account the needs of key partners, and maintaining a relatively balanced allocation approach. For inference needs, we will procure a more diverse mix of chips, including domestic chips from other brands. That is broadly our expected plan.

As for the impact on cloud margins, there are several aspects to consider. First, 2027 and 2028 could be the two years when China's compute supply is tightest, which will support compute pricing and gross margins. At the same time, our proprietary chips and continually expanding scale give us procurement and supply chain advantages. Therefore, compared with using externally procured chips, proprietary chips will improve the margins of our cloud and AI businesses.

Toby Xu:

(Translation)

On your question about T-Head's spin-off and IPO, we recently announced that T-Head has submitted a listing application on a confidential basis. Given listing compliance requirements, we are unable to disclose T-Head's financial information or other non-public information at this stage. What we can confirm is that domestic compute will continue to account for a growing share of Alibaba Cloud's total compute capacity, which is structurally positive for our overall cloud business.

Ronald Keung (Goldman Sachs):

(Original)

So pleasant to see our sequentially narrowed loss for the AI labs and application segment. Just want to ask about we're seeing other China internet giants are actually ramping up on the investment scale on the model and application layer. How does management see the AI value eventually? Would it kind of eventually be at the application agent layer and therefore the investment priority for QwenWork? How do we differentiate QwenWork and the Qwen models and the implications for this segment's investment scale loss for the next few quarters and one or two years?

Eddie Wu:

(Translation)

I think different people in the industry currently have different views on this question. We believe we are still in the early stages of AI. As for whether applications or models will be more important in the future, I actually think the two may need to be looked at together. That is because I believe that, in the future, models and apps will definitely need to develop in tandem, with models helping applications, or applications supporting the development of models. But if an application does not have a stronger model—or one that is both stronger and distinctive—developing jointly with it, I think its long-term user retention is questionable. And if the model behind an application is the same one everyone is using, I think that application will also be highly replaceable.

So our view is that models and applications need to develop jointly in the future in order to create more distinctive product value, or more distinctive product strength. Only then can these products remain in use within enterprises or on customers' desktops over the long term.

Our strategy emphasizes developing our applications and models jointly, to build stronger product capabilities. So, at present, we remain fairly restrained when it comes to more aggressive user acquisition campaigns or user subsidies. Over the long term, we are more optimistic about pairing a distinctive model with a distinctive application to create distinctive value for customers—value that will be relatively difficult to replace. Only this kind of application-and-model combination can deliver long-term value.

Thomas Chong (Jefferies):

(Original)

My question first is about our ARR trend. Given that we have highlighted that ARR exceeded and reached our new milestone, RMB20 billion in August, I just want to get some more color with regard to our full year target. Is it going to beat the USD4.5 billion or RMB30 billion by the end of the fiscal year? On that front, can management comment about the recent trend in ARR and our thoughts about the competitive landscape on Flash model getting increasingly popular?

My second question is about our overseas strategy. Can management comment about our thoughts about the cloud revenue mix between domestic and overseas in coming years?

Eddie Wu:

(Translation)

Overall, MaaS revenue is still fairly dependent on new model launches. In fact, in China over the past two months, the pace of new model launches—including our own and our partners'—

has been somewhat behind original expectations. Our own Qwen4 has also just been released, and at the same time we have high hopes for our next-generation video model, which will be released in the near future. So as we and our partners gradually release new models, I think ARR will see relatively rapid growth in December. So our expectation is that the RMB30 billion ARR target for the full fiscal year remains unchanged.

We are also assessing the impact of the compute export rules on our overseas cloud business, and the specifics will still require some time to analyze and assess. However, at present, Alibaba Cloud operates 109 availability zones across 32 regions worldwide, and data centers including those in the Netherlands, France, and Brazil are gradually coming online. Going forward, the overseas business will continue to be a key focus. For the AI portion, however, we may adopt various cooperation models—for example, we may need more overseas partners—and over the long term, exporting our own chips is also a viable option.

Of course, all of this would come only after domestic demand has been met. But over the long term, Alibaba Cloud has full-stack infrastructure and can serve both Chinese companies going global and overseas customers, and we maintain strong long-term confidence in the prospects of this business and expectations for its development.

Alicia Yap (Citigroup):

(Original)

I wanted to follow up on the cloud margin. Can management elaborate the key growth drivers with the decent improvement of the cloud margin this quarter? How should we be thinking about the cloud margin trends in the coming quarter? Also given the latest constraint on the cloud compute environment, is there a chance that Alibaba might be able to increase the cloud pricings to meet the stronger demand? Would there be any chance on the AI lab losses could also be narrowed at a faster pace amid the latest change of the compute environment?

Eddie Wu:

(Translation)

Several factors have driven the continued improvement in the margins of our AI cloud business. One major factor is that high-margin AI-related products have continued to account for a rising share of our total recurring revenue, which has improved cloud margins overall.

Another factor is the global shortage of compute, which is even more acute in China. In this industry environment, our pricing across products and our bargaining leverage with customers have given us relatively good gross margins. Basically, every product now has a better gross margin than it did last year.

In addition, as our overall data center footprint expands, the economies of scale we generate continue to grow as well, both in the supply chain and on the customer side. Once customers use multiple products of ours, their switching costs increase. All of this, together with customers' continued acceptance of our offerings, is continuing to improve our margins.

Looking at the longer term, as AI technology as a whole develops, the AI cloud business—especially AI cloud services built on our proprietary chips—has undergone a structural change compared with our traditional general-purpose cloud computing business. With proprietary chips, our own models, and full-stack AI services, Alibaba Cloud's offering has in fact already become a non-standardized product. For customers, it is hard to make direct price comparisons the way they would with a standardized product; customers also have differing needs for Alibaba Cloud's full-stack technology.

We have also met with quite a few customers recently. With the arrival of the AI era, Alibaba Cloud is, for customers, not simply a supplier they procure from, but a partner in co-developing products. I think that, over the long term, this is the structural change AI brings to Alibaba Cloud's business.

Therefore, our comprehensive product support for customers, customers' long-term acceptance of our offerings, and the difficulty of direct competitive bidding and price comparison owing to product differentiation—all of these structurally improve the margins of our AI cloud business over the long term.

Lydia Liu:

(Original)

Thank you for joining us today. And we look forward to speaking with all of you again soon.

[END OF TRANSCRIPT]