

Amber International Holding Limited

Second Quarter 2026 Financial Results

September 3, 2026

C O R P O R A T E P A R T I C I P A N T S

Michael Wu, *Chief Executive Officer and Chairman*

Yi Bao, *Chief Operating Officer*

Vicky Wang, *President*

Josephine Ngai, *Co-Chief Financial Officer*

Steve Zhang, *Co-Chief Financial Officer*

P R E S E N T A T I O N

Operator

Good morning and welcome to Amber International Second Quarter 2026 Financial Results.

At this time, all participants are in a listen-only mode. A question-and-answer session will follow the formal presentation. At that time, if you'd like to ask a question, please press star one on your telephone keypad.

As a reminder, this conference is being recorded.

It's now my pleasure to introduce your host, AMBR's official AI agent moderator, MIA. You may begin.

MIA

Good morning and welcome to Amber International Holding Limited second quarter 2026 earnings conference call. I am MIA, AMBR's official AI agent moderator for today's call.

Before we begin, please note that today's discussion will contain forward-looking statements under the Private Securities Litigation Reform Act of 1995. These statements involve risks and uncertainties that could cause actual results to differ materially from those projected. For a more detailed discussion of these risks, please refer to the Company's filings with the U.S. Securities and Exchange Commission, including our most recent annual report on Form 20-F.

Joining us on today's call are Michael Wu, Chairman of the Board and CEO, Vicky Wong, President, Yi Bao, Chief Operating Officer, Josephine Ngai Co-Chief Financial Officer, and Steve Zhang, Co-Chief Financial Officer. Following their remarks, we will open the line for Q&A.

With that, let me now turn the call over to Michael Wu, our Chairman of the Board and CEO.

Michael Wu

Thank you all for joining us.

The second quarter was a strong one for AMBR, and I will start with the numbers. Revenue was \$13.9 million, up 39% from the first quarter. Gross margin expanded to 79.5%. Operating income was \$1 million, and the Adjusted EBITDA was \$1.9 million, both turning positive from losses last quarter. The revenue from our agentic and the digital asset businesses from Amber Premium came in at \$10.1 million, above the \$9 million to \$10 million outlook we gave you last quarter. On our \$50 million repurchase program, we bought back approximately \$2.6 million ADF for about \$5.8 million through June 30, with roughly \$44.2 million remaining.

Those are the results. Today, though, I want to explain why they matter, because they are the first evidence that our new strategy is already taking hold. Two days ago in Hong Kong, we introduced AMBR for what it is today, a Company that builds specialized AI agents. This is a pivot, and I'd rather say that plainly than dress it up as continuity. We built AMBR as a digital Wealth Management business. Now we are becoming a technology company. We're making that choice deliberately and from a position of strength, backed by the numbers I just shared, because we believe this is where the greater opportunity lies.

For me, this is also a return to my original passion and my true ambition. We started in 2017 as Amber AI, and spent the first nine years in markets, learning exactly where capable models stopped being useful to people making consequential decisions. In June, I stepped down as CEO of Amber Group to run AMBR full-time. Building these AI agents is the only thing I plan to work on for the next decade. I say that publicly on Monday, and I will repeat that today to this audience, because you are the ones who can hold me on to it.

The strategy already has two products in market. Ambre is our consumer agent for personal finance. It does what Amber Premium's relationship managers have done for high-net worth clients, portfolio analysis across accounts and asset classes, a daily signal feed filtered to what a user actually holds and actually what matters, monitoring and alerts, and Ambre makes that available far more broadly.

The two design choices matter for this audience. First, Ambre works across users' existing exchange and brokerage accounts. We're not asking anyone to move assets to use it. Second, Ambre does not place orders autonomously, not yet. When a user decides to act, they're handed to our expert team. These were people who were relationship managers, structurers, and traders in our financial services business. We will expand the agent's authority as reliability is demonstrated, not ahead of it. Ambre opened Monday by invitation, beginning with Amber Premium's verified client base.

MIA, your host today, is also our marketing agent, and it's the proof that this model produces real revenue. MIA was built inside iClick and CMRS, our wholly owned marketing businesses, where it runs a substantial share of day-to-day campaign operations for a base of more than 100 enterprise customers. As of Monday, MIA is also available directly, so there's a product, not just embedded in our services anymore. That's the pattern you should expect from us. Operate a business, convert its expertise into an AI agent, and then take that AI agent to market.

Now let's move to the P&L logic, because the repositioning is only credible if the numbers eventually say the same thing as a strategy. Over the \$13.9 million I mentioned, \$7.4 million came from businesses we classified as agentic, including \$3.5 million from A-MM in its first quarter of recognition. That classification reflects how these businesses genuinely run. Operations executed on AI agent structure, not a relabeling of old revenue. Yi will take you through the basics and the details. The reason it matters, agentic is a structurally higher gross margin than the businesses it's replacing, and that makes shift is most of why gross margin reached 79.5% this quarter.

On the rest of the portfolio, as we focus the Company on building AI agents, we're reviewing the shape and ownership of parts of our legacy financial services business. Some of what we operate today fits the strategy as infrastructure. Some might serve clients and Shareholders better in a different structure. We have nothing to announce, and we won't speculate on outcomes, but I'd rather tell you the review exists than have you learn of any outcome code.

I also want to emphasize that Ambre and MIA are the first two AI agents, not the whole portfolio. At our inaugural Investor Day, which we now expect to hold before year end, we'll show you what else we've been building and lay out the financial framework for the transition. What the revenue mix looks like as agentic businesses become the center of AMBR. Until then, our job is simple. Execute on what we just launched, build AI agents, start.

With that, I'll hand over to Yi Bao.

Yi Bao

Thank you, Michael.

The message from my side is simple: the agent strategy is already operating, not just announced. A-MM went from operating system to \$3.5 million of recognized revenue in a single quarter. MIA runs a substantial share of day-to-day campaign operations for more than a hundred enterprise customers. And Ambre is built on workflows our teams have run for years. We build agents on infrastructure we have already proven, which is why we can move quickly without taking on the risk of building from scratch.

What I'd like to do this morning is make the pivot concrete from an operating point of view. Michael described the pattern we follow, which is that we operate a business, we convert its expertise into an agent, and then we take that agent to market. Almost all of the real work happens in that middle step, and it's where I think our advantage sits. So that's where I'll spend my time.

Let me pick up where I left off last quarter. I described A-Suite and A-MM as agent-native operating systems, and a few of you asked afterwards why we were leading with infrastructure rather than with the agents themselves. The short answer is that intelligence on its own doesn't make an agent useful. An agent needs access to the right data, tools it can operate, workflows it can follow, permissions that spell out what it's allowed to touch, and controls around execution, monitoring, risk and compliance. More than any of that, it needs a precise purpose, by which I mean a defined problem, a defined user, and a real situation where the outcome matters to somebody. The model supplies the intelligence, the operating environment lets that intelligence act, and the purpose determines what the action is worth.

That's also why we're not competing at the foundation model layer and don't plan to. We'll stay model-neutral and use whatever intelligence works best for a given job, whether it comes from a proprietary model or an open source one. It's worth thinking about what that means for how you read the model race, because when models get more capable and cheaper, it works in our favor rather than against us. Our costs come down and our agents get better without us spending a dollar on training.

What we intend to own is the layer sitting above the model, which is where you find deep understanding of a particular vertical, the connections into the right tools and data, the design of the workflow and its controls, and the fairly unglamorous work of making general intelligence reliable enough that someone will trust it with a real job.

A-MM is the clearest example of how that plays out. Market making has always run on a fragmented set of workflows, with client requirements sitting in one system, venues and counterparties in another, and then RFQs, contracts, execution, monitoring, settlement and reporting each carrying their own tools and their own manual steps. Our first move wasn't to put an AI interface in front of all that, because an interface on top of a broken process just gives you a faster route to the same bottleneck. We standardized the underlying workflow first, then connected the systems, structured the data, put monitoring and controls around it, and made the process machine-operable one step at a time. That's the layer we described to you last quarter as the A-MM operating system.

In the second quarter, A-MM contributed roughly \$3.5 million of revenue in its first quarter of recognition, and I want to be careful about how I characterize it.

What it tells you is that the infrastructure underneath our agent strategy can already generate economic value. That's a meaningful distinction, because most companies in this field are still asking investors to fund an operating layer that doesn't exist yet. Ours is built, already running, and fits our definition of agentic revenue: the work runs on agent infrastructure rather than through the manual process it replaced.

The part I want to be careful about is what comes next. That revenue today still mostly reflects monetization of the operating platform and the capabilities running on top of it, rather than an agent getting paid directly for delivering an outcome. Our expectation is that the specialized agent gradually becomes the primary interface to that capability, and that the outcome itself becomes what the customer pays for.

Think of the progression as three stages, starting with a manual workflow, then agent-operable infrastructure, and eventually a specialized vertical agent that simply delivers the result. A-MM is in the middle stage today. We'll tell you when it moves, and we'll show you what we measured before we say it moved.

The rest of the portfolio is being built the same way. MIA came out of a working marketing operation, where it learned from real campaigns and real enterprise customers long before we sold it to anyone. Ambre is being built on years of operating experience at Amber Premium, drawing on portfolio analysis, risk monitoring, product evaluation and the accumulated judgment of relationship managers, traders, structurers and product teams. In both cases we started inside an environment we already understood well, converted that operating knowledge into structured workflows and systems, and then let the agent take on more of the work as it earned the right to.

That's also how I'd ask you to think about our legacy businesses inside the new AMBR.

The customer relationships, the domain expertise, the regulatory infrastructure, the execution connectivity, the operating data and the risk and compliance experience all stay valuable, and most of it would be hard for a newer entrant to assemble from scratch. What doesn't have to stay the same is the way we've traditionally delivered those capabilities. A business that grows by adding relationship managers, operations staff and manual steps grows in a fairly straight line with headcount, and that's a different economic shape from the company we intend to build over the next decade. You can already see the difference showing up in this quarter's gross margin.

So as Michael mentioned, we're reviewing where each legacy business and structure fits. Some of those capabilities will end up as infrastructure our agents depend on, and some delivery models will become less central over time.

The pivot is changing how we run AMBR internally as well, because we want to be the first serious users of everything we build. There's a practical reason for that, which is that running our own agents in our own workflows is the cheapest way to find out where they fail before a customer does. It shows us where they come up short, where human judgment is still needed, what context or tooling they're missing, and how the workflow itself should be redesigned. We understand the workflow, we build the agent, we run it ourselves, we fix what breaks, and then we take it outside. That loop is our operating model, and it travels from one vertical to the next.

So when you look at that \$3.5 million from A-MM, I'd ask you to read it as an early proof point rather than a destination. The operating system is the foundation, the specialized agent is the product we're building toward, and the outcome is what a customer should eventually be paying for.

We're early in this transition, but we aren't starting from zero. We're starting with businesses that operate, users who use them, workflows that function and revenue that's already being recognized. The work in front of us now is turning those advantages into specialized agents and scaling the ones that prove they can deliver.

With that, I'll turn it over to Vicky.

Vicky Wang

Thank you, Yi.

Earlier this week, on September 1, we officially unveiled AMBR and introduced the next chapter of our Company, focused on building specialized AI agents for high-value high-stakes. We have been very encouraged by initial response. Since the launch of AMBR, we have seen strong interest from existing clients, prospective users, partners, and the broader market. While we are still at a very early stage, that response has reinforced our conviction that users are looking for something beyond another general. They want intelligence that understands their context, knows what matters to them, and can continuously help them to take action.

This is where we believe AMBR has a differentiated foundation. The AMBR brand is new, but the capabilities behind it have been built over many years. We bring deep domain expertise, trusted financial infrastructure, experience serving sophisticated users, and a detailed understanding of real-world high-stakes workflows. We believe these capabilities become increasingly valuable in an agent AI.

Foundation models are becoming extremely powerful, but in our view, the most valuable specialist agent will require a deeper know-how of the underlying industry. This is where our domain expertise becomes particularly valuable.

Ambre, our flagship Personal Finance agent, is one of the first examples of how AMBR is combining frontier AI capabilities with deep financial expertise to be specialized agents. Over the years, we have built deep capabilities across digital wealth management, risk management, and financial infrastructure. Ambre brings these capabilities together in a much more scalable and intelligent form. Rather than simply providing users with more information, Ambre is designed to understand their financial context, identify what matters most to them, and help turn their intentions into action.

For example, Ambre can build a holistic view of a user's portfolio across different accounts and asset classes, identify concentration and correlation risks, surface the signals that are most relevant to their actual holdings, and continuously monitor specific conditions or tasks on their behalf. We launched the first version of Ambre on September 1 as well, and early response from our existing clients, partners, and broader community.

This is still an early version, and we expect the product to evolve significantly as we validate user behavior and progressively unlock more agent capabilities. Our long-term ambition is that Ambre to make a level of

personalized, always-on professional financial intelligence that historically was only available through a high-touch private banking relationship, accessible to a much broader group of users.

On the other hand, MIA serves the same shape of problems in a completely different market. Marketing teams run research in one tool, insight in another, content in the third, and distribution in the fourth, and nobody owns the seams between them. MIA is built to understand the objective and carry that workflow through end-to-end instead of handling it off four times. We saw it two ways now. Embedded in the services, our marketing businesses deliver and directly as a product, and having both gives us an unusually clear read on what a customer will pay for the agent on its own versus the services rep.

These two markets are picked in the first batch, have almost nothing in common. Personal finance and marketing operations share almost no customers, no regulations, and no workflows. If the same approach works in both, that approach is working, it is not luck. It's also how we will choose the third agent and the fourth. We go where we already operate, where the work is high-stakes and fragmented, and where we hold contacts. A newcomer would need years to assess.

We are still at the beginning of this journey, and there is significant work ahead, but the launch of AMBR marks an important milestone for the Company, and early response we have seen has, again, strengthened our conviction in the direction we are taking. We look forward to sharing more as we expand the capabilities of Ambre and MIA and introduce additional specialized agents across Amber platform.

With that, I will turn it over to Josephine.

Josephine Ngai

Thank you, Vicky, and good morning, everyone.

Before I get into the numbers, let me start with the headline for the quarter. Revenue grew 39% sequentially, and we moved from an operating loss of \$3.2 million in Q1 to operating income of \$1 million in Q2. Importantly, we did that while keeping operating expenses essentially flat at around \$10 million. I think that's an important point for investors. The strategy Michael just described is not being driven by a significant increase in spending. What we are seeing instead is a change in the revenue mix, with higher-margin agented revenue growing alongside continued improvement in our core business.

Typically, when a company goes through this kind of repositioning, you would expect to see a higher cost base first and potentially a need for additional capital. So far, we are seeing the opposite. We are growing revenue, improving margins, and moving into profitability without materially increasing expenses. That's the kind of financial discipline that we want to maintain as we execute this transition.

Let me walk through the quarter in a little more detail. Starting with revenue, total revenue in Q2 was \$13.9 million, up 38.8% from \$10 million in Q1. Beginning this quarter, we have reorganized how we present revenue into two categories, which we think better reflects how the business is evolving. The first is digital assets platform revenue, which was \$6.6 million and includes wealth management, execution, and payment solutions. The second is agentic revenue, which was \$7.4 million and reflects the initial contribution from A-MM, together with our marketing and enterprise solutions business.

Within the digital assets platform, Wealth Management Solutions generated \$5.3 million, compared with \$4.3 million last quarter. That improvement was mainly driven by stronger demand for both our diversified products and several newly launched offerings.

Agentic revenue was one of the key developments this quarter. A-MM contributed \$3.5 million in its first quarter of revenue recognition, while marketing and enterprise solutions contributed \$3.8 million. If you look at digital assets platform together with A-MM, revenue was \$10.1 million. That's slightly above the high end of the \$9 million to \$10 million outlook we previously communicated.

Moving to gross profit, we saw a significant improvement. Gross profit increased to \$11.1 million from \$6.8 million in Q1, and gross margin expanded to 79.5% from 67.7%. The main driver here was the mix of the business. We are seeing a larger contribution from higher margin activities, particularly A-MM and our core wealth management business. From our perspective, it's not just the revenue growth that's encouraging. The quality of that revenue is also improving.

On operating expenses, we remained at approximately \$10 million, essentially flat with the prior quarter. That's particularly important given the growth we delivered during the quarter. We are starting to see the operating leverage we believe can come from deeper AI integrations across business. As a result, operating income was \$1 million for the quarter, compared with an operating loss of \$3.2 million in Q1.

Looking at the bottom line, net income from continuing operations was \$1.5 million, compared with a net loss of \$3.7 million last quarter. Adjusted EBITDA from continuing operations improved to positive \$1.9 million from negative \$3.2 million in Q1, and adjusted net income was \$1.5 million.

Turning briefly to the balance sheet, as of June 30, we have \$34.2 million in cash, cash equivalents, time deposits, and restricted cash.

Let me also address our outlook, because I know this will be an important question for investors. As Michael discussed, the Company is going through a meaningful strategic transition toward becoming an agile AI company. Given that transition, we don't believe our previously issued financial guidance is still the right framework for evaluating the Company's future performance. We have therefore decided to withdraw that guidance while we build more operating history around this field, new business, and get better visibility into their financial contribution. Once we have enough data and forecasting visibility, we intend to provide updated guidance.

I want to emphasize that withdrawing the guidance doesn't change our confidence in the long-term opportunity. It's really about making sure that when we give investors a forward-looking framework, it's based on the business we are building now rather than the business we had before this transition.

Stepping back, I think Q2 gives you more early but meaningful evidence of what this new model can look like. Revenue growth, expenses stayed essentially flat. Gross margin expanded significantly, and we moved from an operating loss to operating profit. For us, that's the pattern we want to continue seeing as we move toward a more agent-native, higher margin business model. We are still early in this transition, and there is a lot of work ahead, but we are encouraged by the progress we are seeing, and we will continue to stay focused on disciplined execution and long-term value creation.

With that, I will turn the call back to MIA. Thank you.

MIA

Thank you, Josephine. That concludes our remarks for today. We will now open the line for Q&A. Operator, please begin.

Operator

Thank you. As a reminder, if you'd like to join the question queue, please press star, one on your telephone keypad. A confirmation tone will indicate your line is in the question queue. You may press star, two if you'd like to remove your question from the queue. For participants using speaker equipment, it may be necessary to pick up your handset before pressing the star key. We'll pause a moment to allow for questions.

Sorry, go ahead.

Michael Wu

Yeah, I see some questions on the screen.

First question, "What is proprietary about your AI agent or about AMBR's AI agent?"

This is actually a really good question, and I'd like to share with the audience our own AMBR's definition of what is even an AI agent. I think the industry sort of comes together to a definition for AI agent. This concept or this whole species is still fairly new in human history. The industry sort of defines AI agents as the model plus the harness. Now, we do believe that definition is incomplete.

Our AMBR's definition for AI agent is an agent is the model plus the harness plus the purpose. I think that will lead us back to the original question. Why do we believe that way? Because the model is the intelligence, and that intelligence is increasing day by day as the model labs compete for better and better models of all sorts, proprietary or open source.

Harness, fairly new concept, essentially means the environment or the tools or setup for that intelligence to do actual work. The best example is coding agents. The harness allows the model to code and write programs for programmers or even non-programmers. We call them biocoders.

But we do think, like our CEO Yi Bao said earlier in his remarks, the purpose is what makes that action from the model valuable to someone. The purpose is essentially what the model is doing for who in what scenario and why. If the model, highly intelligent and increasingly intelligent, acts without the purpose, it's unclear why the customers should pay for that because it's unclear what the value the customer receives.

Now, going back to the original question, what is proprietary about AMBR's AI agents? I do think in the areas we started, Ambre with Wealth Management, MIA with marketing, we understand the purpose, or at least we understand the purpose very well for the customers the original business has been serving for years, for many years. We understand what exactly they need, what are their demands, what are their pain points, and how they like to have these services delivered, how they like to have their problems solved. Now, that is proprietary because any other company with the same model or even the labs that created the model does not have that unless they have done years of servicing these customers like we did.

Take a step back to harness. Nowadays, just like the models, you have increasingly an open-source culture around both the model and the harness. You now have a lot of great open-source models, essentially free to use, free to deploy locally by AMBR or other companies. You now also have a lot of open-source harness. In fact, some of the most popular personal general agent harness like the Permus agent, OpenClaw, Pi Agent, they're all open source, which means anyone, including AMBR, can use them, fine tune them, review them according to our needs.

Now, because we have, again, a very deep understanding about the purpose, Ambre, we understand how these clients like to be serviced around their money, around their wealth management. MIA, we understand how these companies like to be serviced with their marketing. We can then build harness that are special, that are proprietary to these clients, to these purposes, and provide them value in the way they want. In this way they actually find value because the ones who pay, I believe, they find what's valuable.

Now, also, a lot of our harness are also proprietary because they come from the system that's been running, battle-tested, servicing these clients for years. You cannot build these programs. There are a lot of these non-LLM-included, old-school programs. You cannot build these programs. You can't back-code them. They're not battle-tested. They are not the way clients have been serviced or like to be serviced.

Back to the original answer, I think, actually, other than the model, both the harness and the purpose are not only proprietary to the AMBR agents like Ambre and MIA. They are unique with mold that was built over years of servicing these real clients to perhaps one of the highest standards in their industries.

I hope that answers the question. I also hope that provides a bit more insights of how we understand building AI agents given the field is so new. I do think that as a public company doing that, we have sort of a responsibility educating the audience or even potentially sharing what we know about what building AI agents even means.

Can we take questions online?

Operator

Once again, if you'd like to ask a question via the phone, please press star, one on your telephone keypad.

Were there any other web questions?

Michael Wu

We'll take another question from the web. It's a fun one, and I think it can hopefully be insightful for the audience, too.

The question reads, "If investors give your five largest competitors \$50 million tomorrow, what stops them from building Ambre and MIA?"

This is a great question because the answer is both simple and I think, again, comes to how we understand about building AI agents. Frankly, the ones who can build theoretically Ambre and MIA, they do not need that \$50 million. They are potentially the labs or the large internet companies that are already building models and general agents. They do not need that \$50 million tomorrow to build those.

But why are they not building Ambre and MIA? Because that's not their focus. They're fighting different battles. They're trying to build the better model. They're competing very hard. They're the best people. Their billions or hundreds of billions of dollars are spent winning the model war, not winning the vertical agent war we are fighting.

On the other hand, if you give \$50 million to a competitor that wants to build Ambre and MIA tomorrow, it also doesn't help that hypothetical competitor because, first, they are unlikely to understand the purpose we do. Even if they do, they're in the same industry. They are unlikely to have built their harness the way that's agentic or preferable as we have done with A-MM and with the harness around MIA, the harness around Ambre.

Last but not least, I think likely they will be building their version of Ambre and MIA for the wrong purpose. That matters. Likely, they will try to add an AI bot onto whatever they were selling. That's not going to be the right purpose of servicing the users, the customers, like what Ambre and MIA are doing with our customers.

I think this is a great question. I actually think this \$50 million doesn't help any hypothetical competitor of ours because the ones who can do it, they don't need the \$50 million and they're not doing it for financial reasons. They're doing it because they're fighting different battles with their focus, perhaps. The ones who need \$50 million, they probably cannot build what we built or at least in a very extended period of time.

Operator

There are no phone questions at this time.

Michael Wu

Okay. Then we will take one more web question. I don't want to be the only one talking from our team. We have very capable Management. I will start.

“Do you expect to take shares from existing competitors or do you view the segment as open and untapped?”

This is a great question because I think it's a bit of both. It depends on the time horizon, the portion of where our growth or revenue comes from essentially will differ. In the near term, take MIA as an example. MIA is already making revenue from customers. Now, most of these customers probably do not care. They are serviced by MIA or iClick or another marketing company, maybe with or without AI. In that sense, MIA is taking market shares from iClick competitors or even iClick itself. Now, essentially, iClick is—you can think of it as a service arm or an additional layer on top of MIA.

Now, at the same time, I do think customers being serviced directly by MIA are having a very different service experience. They essentially—and especially for a lot of smaller customers, this is essentially the first time they are being serviced and can afford to be serviced by a truly 24/7 complete marketing team. This is a clearly untapped market because these customers, their marketing needs existed before. No way for their marketing needs to be serviced this way didn't exist before.

Now, in the near term, I do think MIA or iClick and the CMRS with MIA behind is operating a lot better than many of its competitors and it will take market share from competitors. But over time, I think the most powerful thing to MIA's service model is it will open up a lot of new customers that they didn't think they can have this level of service.

The same thing happens with Ambre. The level of services Ambre premium team provides to high networks of family, of billion-dollar family of offices was not accessible by smaller investors or individuals in most of the time. It was just too expensive to do that. Now, not only that experience that service can potentially be delivered better, they can be delivered in an affordable cost to a lot more customers. Back to that, before my teammates add more, I think over time it's a bit of both, but a later open market is a lot larger for us.

Operator

Thank you. That concludes today's conference. I'm sorry. That concludes the question-and-answer session. I'll turn the floor back to MIA for final comments.

MIA

Thank you all for joining us today. This quarter marked a clear step in Amber's pivot from a digital wealth management business to a technology company that builds specialized AI agents. Ambre and MIA are the first two agents now in market and the second quarter results are the first evidence that this direction is beginning to show through in the numbers.

We sincerely appreciate your continued trust and support. We look forward to sharing more in the quarters ahead, including at our Investor Day, which we now expect to hold before year end. This concludes today's call. Thank you and have a great day.

Operator

Thank you. This concludes today's conference call. You may disconnect your lines at this time. Thank you for your participation.