BaiRong AI Inc.(6608.HK)

VoxOne - BaiRong Voice Brand Strategy Launch & Investor Salon

Meeting Minutes

Meeting Time: 2026-10-27 16:00:00

[Company Attendees]

Mr. Zhang Shaofeng - Founder, Chairman of the Board, CEO

Mr. Ji Wenyao - Head of AICC (AI Contact Center) Business

Mr. Starry Chen - Head of R&D

Ms. Yang Ying - Head of Finance

Ms. Kang Yicong - Investors Relation Director (IRD)

[Transcript]

BaiRong, IRD

First, welcome everyone to BaiRong's VoxOne launch event today. Next, we will move into the investor Q&A session. Many of you are probably familiar with me - I am Kang Yicong, the company's Head of Investor Relations. Thank you again for staying after the launch event for this closed-door discussion. The VoxOne voice technology brand has just been launched, and I'm sure you have many questions, so we will dedicate most of this session to Q&A. Before we officially begin, let me note that our founder, Mr. Zhang Shaofeng, is currently being interviewed by media and will join us a bit later. During the first part, our AICC (AI Contact Center) business, R&D, and finance leads will engage with you, and Mr. Zhang will join later. We appreciate your understanding.

Let me introduce the three management members present today. First is Mr. Ji Wenyao, Head

of AICC (AI Contact Center) Business, primarily responsible for the commercialization and implementation of AICC (AI Contact Center) business. He will address your questions regarding AICC (AI Contact Center) business expansion and delivery. Second is Mr. Starry Chen, Head of R&D, responsible for the development of foundational voice models and how Agent technology is integrated into our products. You're welcome to ask him technical questions about the BaiYu full-duplex voice interaction platform launched today. Next is Ms. Yang Ying, Head of Finance, who will address BaiRong's finance-related questions.

For today's closed-door session, if questions involve undisclosed material information, we may not be able to respond. Please also note the investment risks associated with forward-looking statements. Let us now officially begin the Q&A session. Please raise your hand to ask questions. Let's start with the first person in the second row.

Analyst

Thank you for the invitation. I'm an analyst from the Emerging Industries team at Tianfeng Securities. I have a few questions regarding the AICC (AI Contact Center) business. The first question is: what is the latest progress on AICC in new scenario business? During the launch event, a partnership with a logistics industry client was mentioned. Besides this logistics client, what is the status of engagement and deployment with other logistics clients? Also, beyond logistics scenarios, what is the progress in other scenarios? That's my first question, thank you.

BaiRong, Head of AICC (AI Contact Center) Business

Let me answer your first question. AICC's new scenario revenue for the first half of the year has been disclosed in the public financial report, at RMB 11.26 million, with a year-on-year growth of 195%. As you can see from the report, the number of clients increased from 40 last year to 140. The core driver remains the benchmark client in the logistics industry. We started in June and moved quickly - by July we reached approximately 15,000 and maintained that level. As the client mentioned, and this is a forward-looking statement, we will soon transition to a larger volume because we are already testing full-scenario deployment. A portion of full-scenario traffic has already been diverted, and the actual volume already exceeds 15,000.

However, full-scenario deployment is still in the production verification phase, and the latter

part has no financial contribution yet, so I cannot share details. If all goes well, large-scale traffic diversion will begin by the end of November; if progress is slower, it should be completed around December. Regarding other clients, you can see from the clients present today: all major express delivery groups are here, except for one that couldn't make it.

The logistics industry is actually divided into two categories: individual parcels and bulk freight. We are basically starting with large clients, following the logic of first conducting POC with a sample client, expanding from a single scenario to full-scenario deployment, and after becoming a benchmark client, replicating across the industry. So our current path continues to advance along our overall business objectives.

The second area is the securities broker sector, where we are currently in the contracting and implementation process, with 13 firms. The deployment scale of silicon-based employees at a single top-tier securities broker continues to grow, with approximately 60+ scenarios. Airlines are currently in the production traffic diversion preparation phase - I cannot disclose which specific airline. After airlines, our second priority is followed by banks and telecom operators. All the clients present today are clients we are currently implementing or operating with. That's the overall situation.

Analyst

Thank you. A follow-up question - regarding the processing volume you mentioned for the logistics industry client, what is the unit?

BaiRong , Head of AICC (AI Contact Center) Business

Inbound calls.

BaiRong, IRD

That is, the number of incoming calls.

Analyst

Understood. Second question: you mentioned starting with logistics and now expanding to airlines, telecom operators, and securities brokers for cross-industry expansion. What is our cross-

industry strategy, or what are the main reasons we can successfully replicate? Also, facing state-owned enterprises, financial institutions, and SMEs, these clients have very different requirements. How do we balance customer acquisition and delivery revenue?

BaiRong, Head of AICC (AI Contact Center) Business

Let me answer your question. Everything announced today revolves around a voice foundation platform. The voice foundation is the basis of all productivity, including how to respond to customers and how to answer - this is a universal capability. What's missing? Industry-specific knowledge - how each industry should respond, which requires building a vertical domain model. This ensures the entire tech stack is universal, while the business stack is delivered through FDE for the last mile to ensure successful outcomes.

Technically, please don't think this is still traditional ASR and TTS. At today's conference, we discussed vertical domain large models. We establish independent vertical domain models for each industry, adding corresponding knowledge points on top of general interaction capabilities to address the delivery path for state-owned enterprises and SMEs. Of course, each industry is different. We primarily promote cloud deployment; the second option is hybrid cloud deployment, where customer information stays within their domain while leveraging our computing power; the third is the LITE and AURA products launched today, using different packaging formats for different scenarios. However, for our technology foundation, it's the same - just packaged into different scenario agents. The integration with client systems is unavoidable for any enterprise. This reduces the overall cost of usage for enterprises, as clients always make quick decisions based on whether it's cost-effective. I hope this answers your question.

BaiRong, IRD

Let me interrupt the Q&A for a moment. We will now officially welcome our founder, Mr. Zhang Shaofeng, to join the discussion.

Analyst

Thank you very much. I have two more questions. The third question is about the RAAS model proposed at the launch event, with three pricing approaches: per AI-completed task, per effective

call or agent headcount, and revenue sharing. How are client-side acceptance metrics structured in contracts for pricing? For example, how do we view and bear risks when results fall short, timelines are long, or call volumes fluctuate? Additionally, what direction does the company hope to guide the pricing model toward ultimately?

BaiRong, Head of AICC (AI Contact Center) Business

From the overall business structure perspective, we ultimately aim for outcome-based pricing. But some things are measurable and some are not. In the customer service domain, whether for airlines or logistics, we promote per-task pricing - if a call is transferred to a human agent, that task is considered failed. This is task-based. Per-agent pricing applies when services are difficult to measure by outcome - they're not task-based but rather consultation providing emotional value, as commonly seen with airlines: How is my flight today? This falls into two categories: one is ticket changes/refunds, which is easily completed - change confirmed, SMS sent, customer pays, that's a task; the other is consultation, such as where to find luggage, which is hard to measure by task metrics and may be priced per agent per day based on call volume. So there are roughly two approaches: top-tier clients are always priced by measurable tasks, the second by agent headcount. Additionally, there's revenue sharing, which is essentially pricing based on business outcomes.

BaiRong, Founder, Chairman of the Board, CEO

In practice, there are demand fluctuations. For example, during Double 11, I might have only 10 orders for a campaign, hypothetically speaking. If it surges to 30, this kind of elastic demand fluctuation places higher demands on operations management. So the problem AI solves is not just cost - it's the flexibility, elastic scaling, and training costs I mentioned.

BaiRong, Head of AICC (AI Contact Center) Business

To add to what the Chairman said, this is essentially about elastic costs. This has been a long-standing issue in the industry - in my past work experience, I managed a large-scale operations team. This kind of elastic staffing isn't about 1x or 2x - it's 5x or 10x. Take a simple example: investors here frequently fly, and when you encounter thunderstorms and want to change your

ticket, try calling customer service. The airline industry's average answer rate is the worst - only about 70%, and during peak times as low as 40% or even lower. So the problem they need to solve is ensuring every customer gets through. When we spoke with an airline, they explicitly asked whether we could handle all these incoming calls, because there are regulatory compliance metrics for airline management. They felt it wasn't about how much money to spend, but about meeting compliance requirements for improving this metric.

Analyst

My last question. This year, our company's progress in AICC has been very rapid, both in commercialization and subsequent client contracting. From a technical perspective, what are the differences between the new generation of AICC and the previous generation of technology?

BaiRong, IRD

I'll ask our Head of Business and Head of R&D to answer this together.

BaiRong, Head of R&D

The AICC or human customer service / intelligent customer service scenario has existed for a long time - over ten years. The earliest approach was based on regex keywords, and later when BERT models emerged and semantic classification matured, people used semantic intent classification models. But honestly, the problems we now solve with AICC are all problems that previous solutions couldn't solve - because tasks that previous solutions could handle wouldn't end up at human agent seats. What we need to solve today are the difficult problems left behind by the previous generation. For example, many demos at today's event, like multi-intent processing: a user simultaneously requests two things, but previous intent-based solutions could only handle one and would miss the other.

The second is tool calling. The old NLP approach might just do one round of slot filling, extract relevant fields, and call a tool. But if the tool call fails, how do we switch to a new tool or find an alternative path to solve the problem? This is the significance of AICC today - because this agent has autonomy. We tell it the goal, and it tries different approaches on its own to accomplish the task. This is completely impossible in traditional intent classification-based models.

BaiRong, Head of AICC (AI Contact Center) Business

Let me explain from a business perspective why I chose the customer service track. It's precisely because of what was just said: many times, the previous generation of technology simply couldn't solve the problems. For outbound calls, it's a fixed process where I constrain the customer within a specific workflow; but for inbound calls, you don't know what the customer is calling about - whether it's to track a package, complain about a delivery person, or something else entirely. In this unknown situation, with traditional NLP keyword approaches, if the customer says I want to complain, they'd be transferred to a human agent. Why do they want to complain?

Because the system can't align with them like a human would. When a customer says I want to complain, our AI first asks What seems to be the problem? - this creates a pressure buffer zone that human customer service would normally handle. I listened to 20 call recordings two days ago, and many were like this: the delivery person hasn't contacted me for three days, you promised delivery by tomorrow, you keep saying you'll contact me, why haven't you? Then we need to initiate follow-up actions, including proactively calling those delivery personnel. This is not simple ticket processing, because in the industry, customer service is a starting point, not an endpoint - this is a core issue. Traditional NLP has no idea who to contact downstream to resolve the issue.

Additionally, in terms of the overall scenario, traditional NLP can only handle single-turn problems and can't handle long memory. If I tell it a lot of things, it just follows the script saying try again, and if I suddenly want to go back to a previous issue, it can't - because it's a tree-structured flow, unless the branching point allows going back, similar to neural networks. When I was doing NLP two or three years ago, I kept thinking about solving this. Why didn't we touch customer service two years ago? Because the technology hadn't broken through yet. This year, we've achieved a complete breakthrough and found the right clients. We're moving forward.

BaiRong, Founder, Chairman of the Board, CEO

Let me add - there are three major differences from the previous generation of technology.

First, can you truly understand human speech completely? The previous generation of NLP simply broke a sentence into individual keywords, essentially doing keyword matching: splitting the user's words into tokens and checking if they match your lexicon, then performing actions based

on matches. So it couldn't understand the user's meaning from the complete sentence. This is the first major experience improvement - genuinely understanding human speech, which is very easy for large language models.

Second, after understanding, can you speak like a human? Previously, when you called intelligent customer service, you didn't want to continue the conversation as soon as it spoke, because the voice wasn't realistic enough, was very mechanical, and the grammar was either too formal - too formal is actually not good; it should be casual, colloquial, even with some filler words. Can the speech not be standard Mandarin? We found that the best results aren't with standard Mandarin but with Mandarin that has regional accents, because people from different regions all speak with accents.

The third key difference is whether you can invoke operational systems to complete orders, modify tickets - that is, tool calling capability, previously called function call. A human agent can open a browser, log into the system, and modify an order, but traditional NLP cannot. If you can't do the third point, won't you end up transferring to a human anyway? So whether you can write back to operational systems - that's the domain of traditional software. Can I write an order back or modify it?

These three major differences: first, genuinely understanding human speech; second, speaking like a human; third, getting the job done. My summary today is: listening, speaking, thinking, and acting simultaneously - the thinking layer is invisible to users. Take an example of a fuzzy requirement: we're currently testing with an airline. A customer says they want to fly to a destination with a layover of no more than 3 hours. How would this traditionally be handled? A human would log into the system, search, find two connecting flights, and review them manually. Traditional solutions can't do this - you need to search, so that's thinking. Finally, placing the order and confirming it requires modifying the operational system - writing back, which requires write permissions. Previously it was read-only; write permissions are very strict, requiring database modifications - either inserting a new record or modifying an existing one.

Technically speaking, these are the three major differences: understanding human speech, speaking like a human, and being able to invoke and write to operational systems. Previously, all

three basically relied on human agents. So when people talk about intelligent customer service, the term has been overused, and people feel this field isn't exciting - many such companies are struggling. Why? Because enterprises didn't feel it solved their core problems - human involvement was still needed. So what clients say today is very practical. At one of our logistics industry clients, our competitors are exclusively major tech companies - no small companies can afford to play this game. Because BaiRong's underlying models are self-trained, other application companies aren't willing to invest heavily in large models, while major tech companies can't focus on vertical domains.

This isn't purely a model problem either. For example, when communication volume increases, there are communication congestion issues, and we need to solve communication soft-switch problems - that has nothing to do with AI. But if you don't solve it, the overall performance suffers. So clients say: no one dares to commit to outcome-based pricing, because there are intermediate steps that aren't under their control. If performance is poor, the other party says it's not their fault and won't pay, which is unacceptable. But BaiRong actually took this on - it's tedious work that requires years of accumulation.

I think for this type of client, it's very simple - look at the transfer-to-human rate. The ultimate metric is: if transferred to a human, we don't charge. Clients ask directly - do you really dare to commit to not charging for human transfers? We are the only ones who do this.

Investor

Right, it's about the client's IT capability, or their ability to integrate with our product and system - they need to be prepared to follow this approach.

BaiRong, Founder, Chairman of the Board, CEO

It's about whether your AI can do function calls and invoke their systems. The best case is when systems expose APIs; if not, it's difficult - we need to use computer use to log into their software system, simulating human operation, which is slower.

Investor

Reading from the client's system should be relatively easy, but writing to their system requires

them to be prepared.

BaiRong, Founder, Chairman of the Board, CEO

When you say prepared, do you mean prepared to use AI, or that their systems can be opened up? System openness isn't our concern - traditional software can do that, and theoretically everything has an API. Any mature software company - Oracle, ERP, Yonyou, Kingdee - all have APIs, because they need to integrate with other office software, with DingTalk, Feishu, etc. It's not difficult.

Investor

The crux of this question is: what will be the pace at which clients across industries adopt our solution?

BaiRong, Founder, Chairman of the Board, CEO

Basically our approach is: first, use our own resources to target industries with high call volumes, such as logistics, airlines, and telecom operators, to establish benchmarks. The goal is to capture three of the top five players - if three of the top five are your clients, you've established the standard. Then we use lightweight versions, like BaiRong LITE and AURA, distributed through channel partners and agents.

Investor

What's the route?

BaiRong, Head of AICC (AI Contact Center) Business

Let me add to what the Chairman said. I understand your question is about API interfaces or whether clients are prepared. Most clients probably aren't prepared, but during the engagement process, once they see from our actual deployments that this helps them - including airlines - they gradually open up. They don't wait for a specific moment; because they haven't encountered a company like this before. Why would they need to open up if no such solution existed? When they see this solution exists, they realize they need to open up to complete the functionality, and they proactively come to you. Let's say I need to navigate 5 interfaces to execute a task, connecting 4

backend systems - I know those 4 systems' interfaces because the original real-time text interaction interfaces already exist; it's just whether they're willing to open them at this point. From my experience, they are willing - as long as they accept this logic, they'll open up.

BaiRong, Founder, Chairman of the Board, CEO

I think some industries are special - for example, e-commerce platforms' last mile is mostly franchisees and agents, with few headquarters staff and demand at the local outlets. It actually depends on how much of the last mile we can handle - as our share increases, business revenue is expected to grow accordingly. It's actually divided into inbound and outbound - this industry has a large inbound volume, but the outbound notification and marketing volume is even larger.

Investor

But that's not our strength.

BaiRong, Founder, Chairman of the Board, CEO

We do that too, but the competition is more intense, because traditional outbound calls are brief notifications - maybe three or four rounds and done.

BaiRong, Founder, Chairman of the Board, CEO

Our current strategy is to first handle high-value tasks, then the rest. Clients are already starting to use us for outbound calls too. Initially we didn't want to do it - I said it's too cheap; but then the client said they could accept a slightly higher price because the results are better. So gradually, if competitors exit the outbound market, we can take that over too.

Investor

So it's like our model portfolio: high-value, high-intelligence models for outbound, and another model for inbound. For the latter, we use a lightweight model configuration with minimal computing power consumption and a different pricing strategy, bringing that business in too. That wouldn't be outcome-based pricing but volume-based pricing.

BaiRong，IRD

Great, thank you. Let's invite the first investor in the third row.

Analyst

Thank you for the opportunity to ask questions. I'm an analyst from Huatai Securities. I have some questions about our AICC (AI Contact Center) business progress and financials. First, regarding AICC (AI Contact Center) business metrics - are there any updated data you can share with investors, including client numbers and business progress? That's my first question.

BaiRong, IRD

I'll ask our Head of Business and Head of Finance to respond.

BaiRong, Head of AICC (AI Contact Center) Business

AICC is currently in a new business development phase. We don't disclose segment-level details - we can only share an overall figure, not individual client data. But as you can see from the first question, we are continuously expanding into new scenarios and industries. Our client base is quite concentrated - logically, any large enterprise scenario with massive C-end interactions is our client. For example, logistics, airlines, banks, securities brokers, and telecom operators - these are our five key industries, and there are no entirely new industries yet. From a metrics perspective, any industry with over 100 staff and existing contact centers, we will assess its SAM. These top five tracks are already large enough, and we will continue advancing with this strategy.

Based on industry research, we have already begun full-process traffic diversion. A magnitude-level increase for a single client doesn't mean our overall revenue will rise quickly. It will definitely require a cycle.

For the airline industry, we're in production verification. Securities brokers and banks are also being pursued. Banks are more complex because large language models now require regulatory filing, with the entire filing cycle taking 3-6 months. This is also why today we invited a senior executive from a leading financial industry institution - we've reached the final stage, which is deploying computing infrastructure within the institution and have begun serving a small batch of customers. However, these figures are not convenient to disclose at this stage.

Analyst

Thank you. I have another financial question that many investors are concerned about. Our

first-half profitability is under some pressure, and I'd like to ask management for their outlook on future revenue and profitability. Also, people are interested in the revenue growth of the new AICC (AI Contact Center) business, as well as the traditional AI Decisioning Business and AICC for the Lending Use Case. What's management's view on the future? Thank you.

BaiRong, Head of Finance

Thank you for the question. For our business this year, the first-half decline is visible and was passive, driven by changes in the external environment. Looking at the second half, the external environment hasn't changed. But this industry change occurred last year, so we've taken many management actions this year. The cost savings from these management actions aren't clearly reflected in the first-half financials because they take time, and will be reflected in the second half. Since the revenue decline trend hasn't changed, we will continue these management actions in the second half, which will also affect full-year costs.

But for next year, we'll be lighter on our feet with continued management cost reduction. I can't share specific numbers, but I can say that with revenue continuing to decline in the second half, thanks to effective cost control, the second-half loss will definitely narrow. Investors who know BaiRong well should also know that BaiRong's operating costs mainly come from personnel and R&D investments, which are substantial. Since the AI Decisioning and AICC for the Lending Use Case lines are both passively affected by industry changes, the corresponding costs will decrease, and we have corresponding actions every quarter. As for the AICC for New Use Cases growth, as mentioned earlier, volume will ramp up around November or December. We maintained over 190% revenue growth in the first half, and this is still being maintained in the second half even without the new volume ramping. Next year, as the business advances, there will be a more specific cadence worth watching. That's roughly the situation.

Analyst

Thank you.

BaiRong, IRD

Great, let's invite the third person in the second row.

Analyst

Hello everyone, I'm from Orient Securities. I'd like to ask about the customization cycle for expanding into leading clients across industries. How long is the typical cycle for industry-specific and client-specific customization? And at what revenue scale from leading clients might we see meaningful business progress? That's the first question. The second question is about the technology leadership - you mentioned significant technical advantages. How do you assess the difficulty and timeline for competitors trying to catch up? Those are my two main questions.

BaiRong IRD

Let me ask our Head of Business to answer this first.

BaiRong, Head of AICC (AI Contact Center) Business

Let me answer the first question about timelines. The airline and logistics clients we've targeted actually have relatively high standardization levels. For the first client, the timeline was honestly quite long. Why? Two reasons. First, our FDE team needs to first understand their business, then perform engineering decomposition and alignment. The core is understanding their business structure - how to build the agent, what business problems we need to solve.

For example, with a securities broker we're currently working with, they want the entire AI to become a super agent that can answer any question regardless of who's asking. But during implementation, we don't do it that way. The reality is: what are the 60% of issues that come in first? We solve those first, then tackle the next 30% and 20%, using this path to ensure rapid revenue deployment. Otherwise, if you demand a super agent from day one, you can't even go live for a 1% business - I think that's not great for either party. So we process it in this cadence. As for the ultimate ceiling, from the agent perspective, the capabilities have already been accumulated.

Where's the difference? The difference is in the client engagement process: what interfaces they need to open, what business processes are involved - there's some customization. But for VoxOne itself and the agent construction, the changes are very minimal. At least across the projects I'm implementing, I've seen this trend. For example, package tracking works the same way at one

express company as another - they just open interfaces differently or have different system concurrency levels, but the logic is the same, given the client's level of openness. So between the first and second client, at least in the FDE phase, the time will be significantly shortened.

Looking at the medium to long term, once templates are formed, even with LITE all-in-one machines, I can just deploy a unit - you want 10, I give you 10; you want 20, I give you 20. This is our basic judgment. Look at major airlines - behind them is just one system provider. The largest system providers behind logistics companies are also just a few. It's the same in the US market - behind major businesses are always major vendors. Like Huawei domestically - once I finish integration with them, everything else becomes simple. That's roughly the picture.

BaiRong, Founder, Chairman of the Board, CEO

Let me add. Our capabilities can be split into two categories: one is industry-agnostic, such as the large language model itself, communications engineering, and soft switching - these are unrelated to specific industries. What's industry-specific is the industry's products, services, and scripts - each industry is different. So for the first client, most time is spent understanding this. Of course, setting aside cases where the client can't get their act together or has complex internal approval processes - those we can't control. What actually takes our time is: understanding the first client in a new industry - how this industry works, what products and services exist, how to respond to issues, and getting their historical script library. This takes time. But once we've nailed one client - say one logistics client - moving to another, we can basically go live quickly. The technology part, as I just mentioned, is fully reusable. What's not reusable is the business scripts for different industries.

BaiRong, IRD

Regarding your follow-up question about the scaling progress of this business, let me add some context. Although I can't share specific numbers, I can provide a business development framework for your reference. Whether you look at our AICC for New Use Cases or what we call AICC for the Lending Use Case, essentially it's the same AICC technology applied across different businesses and industries - it follows industry lines. For BaiRong, going back to before 2025, BaiRong had already achieved scaled revenue in the credit industry using AICC technology.

This business line, once it achieves scale effects, has promising business prospects - you can reference and evaluate this independently. You're also familiar with this business line's process - it requires some customer acquisition investment, so there will be sales expenses. But even so, once scaled, it can achieve what we consider a healthy business development level.

Returning to AICC for New Use Cases. As you just came from the main hall, we've been emphasizing inbound calls. What does inbound mean? It means when business occurs, from a cost structure perspective, I don't need to increase sales expenses. Looking at the overall cost structure, because it requires large models, there may be some computing cost increases as scale grows. Of course, we'll employ various technologies, including FDE teams and Harness engineering optimizations, to ensure computing costs don't cause significant cost fluctuations.

From this perspective, with an excellent business structure and ideal cost structure, the ultimate business performance will be better than our previous AICC for the Lending Use Case - this is our expectation. Previously, AICC for the Lending Use Case reached roughly RMB 900 million, or even over RMB 1 billion in scale. So you can think back to why we emphasize the SAM of over RMB 170 billion: in the medium term, say 3-5 years, if BaiRong captures just 1%, 2%, or 3% of this SAM, the impact on BaiRong's overall revenue would be a very significant change, or a very substantial growth expectation.

BaiRong, Founder, Chairman of the Board, CEO

Let me add. Not only is the deployment model cloud-based, but the client doesn't even have their own operations staff - everything is fully outsourced to BaiRong. The downside: the upfront costs are indeed significant. The upside is that we've become domain experts, and replacing us or competing with us becomes difficult, because you need to understand not just technology but also the industry. BaiRong previously achieved this in the credit industry. In a normal operating environment, the business growth we saw last year was extremely rapid. Because the more steps you entrust to us, the more revenue we generate from each step - earning multiple times across multiple steps. The upfront investment is all loss, but the barriers become very deep, except for external environment risks that require ongoing attention.

The downside is the high upfront cost. The first year's investment has been made; the second

year may not require additional investment. But with the BOX or even phone-form products launched today, which are highly standardized - a software-hardware integrated solution where you have everything in hand and don't really need our operations - a small enterprise can use it directly, and we would definitely be profitable, otherwise we wouldn't sell it.

For a logistics industry client like this, it's not easy for major tech companies to compete with us. They still charge per call or per minute, not outcome-based. This has pros and cons. But look at our credit industry example - we went deep, and once mature, the business performance was actually quite good.

BaiRong, Founder, Chairman of the Board, CEO

This is the most financially difficult period in BaiRong's history, because we're undergoing a transformation. In other industries, you have to go through the same path the credit industry went through. But we're fortunate - a logistics industry client case has proven that the full-service model works, still on an outcome-based pricing basis. Previously, people questioned us: you do outcome-based pricing in financial services, can you do it in other industries? We couldn't answer. Now at least one logistics client has answered this question.

BaiRong, IRD

Great, thank you. Regarding how to maintain technological leadership? Let me first ask our Head of R&D to answer, with any additions.

BaiRong, Head of R&D

I think this breaks down into several layers. First, for solving industry problems, it's about industry understanding. Our underlying tools and technology layer are unified, but the FDE and industry expert work on top - if major tech companies don't invest time and resources here, it's hard to improve.

Next is our voice foundation platform, the Harness layer. Today you can see many Harnesses on the market, but they're all designed for text. Ours is different - specifically designed for voice scenarios. Voice scenarios involve many different tasks from text. For example, we demonstrated this afternoon: how to handle noise suppression, multi-speaker scenarios, distinguishing effective

semantic content from noise in user speech; or when extracting information like ID numbers or phone numbers from speech, how to determine completeness and what needs supplementing or correcting. All of this is packaged in our voice Harness layer. In China, whether you look at open-source or closed-source products, very few do this. Why? Because the technical threshold is higher than text. We haven't seen a mature platform in this space either - our Harness platform layer has its uniqueness.

Below that is the model layer, where we have the deepest accumulation. Over the past ten-plus years, we've accumulated massive data in telephone voice scenarios. As the two investors were discussing, our models are self-trained - both pre-trained and fine-tuned. We have multiple sources: some based on continued pre-training of open-source models, some based on fine-tuning open-source models, but we emphasize our fully self-pre-trained models.

For example, voice scenarios have different sampling rates. What does that mean? Voice scenarios involve converting analog signals to digital - how many samples per second. In our telephone scenario, most audio is 8K sampling rate; but in the voice systems of major tech companies, sampling rates start at 16K, and much real-time voice communication uses 44K sampling rates. This makes their models more delicate, dependent on large volumes of high-quality data. In our telephone scenario, we only get 8K sampling rate - sparse data with fewer features, so some models simply can't perform well in this environment.

On this front, there were also some announcements at today's event. We've open-sourced some of our models for researchers worldwide to evaluate in real-world scenarios. On the VoiceBench leaderboard, we achieved first globally. Major domestic and international labs and frontier research labs are all on it. VoiceBench is the world's first benchmark for evaluating Agent capabilities in complex voice scenarios. As I mentioned, in highly complex audio environments with user noise and multi-speaker situations, our processing capability is ranked first globally.

Additionally, when users speak, there are pauses - which are effective interruptions versus invalid input; various scenarios also have jargon and terminology that standard ASR may not recognize, and we have solutions for these too. We have duplex detection models that feed audio to the agent to determine whether someone has finished speaking and whether the utterance is

meaningful. Our capability in this model is also first globally, and we've open-sourced it for developers worldwide to test.

To summarize, delivering a good voice agent requires three layers: FDE industry experience, a strong orchestration system, and very strong model capabilities. These constitute our technical barriers. For a competitor to enter, I think it would take considerable time, because we didn't start building these tools last year - we've been continuously investing in this area for about 5 years. This is the result of continuous accumulation in engineering, data, and algorithms.

BaiRong, Founder, Chairman of the Board, CEO

Let me add one point. If you don't follow this closely or aren't particularly familiar with it, you might not know there's a key issue here. When you talk to a general-purpose voice assistant, it feels very real, right? Some companies previously used a general-purpose voice assistant, but later switched to BaiRong. On the internet, when you make a WeChat call, the sampling rate is at least 16,000 times per second - some are 32,000, some 44,000. Real telephone calls only sample 8,000 times per second. This means internet calls have high information density and better data quality. It's like an athlete who trains drinking Red Bull but can only drink mineral water during the actual fight. So if you take a model trained on 16K or 32K and move it directly to 8K, performance drops sharply. Our original model is trained on 8K, because we started from the streets. Our 2.6 billion calls are all real telephone channel data - only 8,000 samples per second, so it's much harder. Only 8,000 data points per second; while internet calls and general-purpose voice assistants use at least 16,000. If you're not in this field, you wouldn't even know this difference exists. So once you move a general-purpose voice assistant model to telephone, the performance immediately deteriorates. And as I mentioned, it also involves communications, which has nothing to do with AI. You need to solve many problems unrelated to AI to make it work end-to-end. If you don't solve them, performance suffers, so no one wants to commit to outcome-based pricing - only BaiRong dares to do this. There are many problems that have nothing to do with AI that we had to solve.

BaiRong, IRD

Great, thank you. I see an investor in the back has had their hand up for a while.

Investor

Hello, I'm from Junshan Capital. Building on the financial question just discussed - regarding the negative impact of external environment changes on our overall revenue, does management have a rough estimate of how many more quarters this will continue? Thank you.

BaiRong, Founder, Chairman of the Board, CEO

First, this won't get worse. Month over month, it won't be worse than now; year over year, whether next year will decrease further is uncertain. Because the massive impact really started after April this year. Actually, we started recovering in March-April and thought we could bounce back, but then some unexpected events occurred. But month over month, it won't be worse than now. Year over year, it might still worsen, because Q1 this year was still okay.

Investor

We haven't deliberately cut this part of the business.

BaiRong , Founder, Chairman of the Board, CEO

That's right - our own energy and resources may shift toward other areas, certainly. But I wouldn't refuse orders from clients - we haven't done that.

BaiRong, IRD

Regarding resource allocation for this business, as the Head of Finance explained, we've already made corresponding reductions, and will monitor the situation going forward.

Investor

Okay, thank you.

BaiRong, Head of Finance

Let me add a few points. We're also working on next year's budget and can see recent trends. Two parts: one is AICC for the Lending Use Case, which is directly tied to transaction volume and has been the most significantly impacted by external environment changes. If you regularly follow our financials or industry trends, it has dropped rapidly to a relatively low level. While it will continue to decline next year, the rate of decline will decrease given the already low base. The other part is our AI Decisioning Business, which isn't tied to transaction volume but provides

services. We're already seeing a stabilization in this segment - a stabilization that feels more solid than the first half. It seems to have genuinely stabilized.

BaiRong, IRD

Let me add a point. Our current focus, including today's VoxOne launch event, signals that the company's future strategic focus and resource allocation is on AICC and new scenarios. Earlier, investors asked many questions about client expansion and our pipeline clients. If you noticed, many of today's first-row guests are from the aviation industry, leading financial institutions, and the logistics industry. The fact that they're willing to attend and some are willing to speak on stage is because of our partnership and their recognition of our products and technology. Commercial discussions are also ongoing. So while I can't make specific guarantees, from this point in time, financially BaiRong may be in the darkness before dawn; but from a business and strategic perspective, we're very clear that AICC is our direction. On the business side, we see many clients already in deep discussions with us - not initial contacts, but progressing toward POC. For the company's next 1-2 years, our morale is very positive about this.

BaiRong, Founder, Chairman of the Board, CEO

To summarize, this is indeed the most difficult financial period since the company's founding, because the existing business decline is severe, and revenue from organic growth can't keep up with the decline caused by external environment changes. But the trend we're seeing is that we've successfully replicated the outcome-based pricing model from the financial industry to non-financial industries - like express logistics, which is simpler; and now we're moving to the second industry - airlines. An airline has passed internal testing, and the next step is to put it on real traffic with real customers using it, starting with a portion diverted to us. We might not even charge initially, but they've agreed to the outcome-based pricing model. So I think this business model - over the past 1-2 years, some investors worried we couldn't replicate it, but in fact, as long as it's outcome-based pricing, basically no client refuses, because they only partner with you when they're making money. If we sell boxes, that's different - I might just give you one and you figure it out; but for the full-service model, we commit to outcome-based pricing. It's not easy for others to replace us.

Regarding competition - competition isn't purely about technology. It's also about understanding the industry and whether you're willing to solve problems unrelated to AI, like communications. Major internet companies have very strong internal technology, but they simply assume their clients are equally strong, which isn't the case. Once they try to replicate it, the client's surrounding infrastructure isn't adequate, so performance suffers. Many major tech companies are like this because they haven't done B2B - they simply assume their client's environment is as good as theirs, when it's actually much worse. BaiRong learned this the hard way. I know clients aren't like that - you need to think through everything for them so they can achieve the same results as you. So the capability building here - first, it's not purely technology; second, the technology isn't purely about the AI model alone. Our chain is very long, with perhaps 5-6 key links. If any one breaks, the final result suffers.

BaiRong, IRD

Great, thank you. Due to time constraints, we've already been going for 1 hour and 15 minutes. I'll take one final question - anyone? Let's go with the first person in the first row.

Analyst

This question may be more long-term and organizational. You specifically emphasized the AI native organization, and in this morning's panel discussion, you made a point that was very inspiring to me: you said AI is better at combining multiple different job functions, like sales and customer service. We see that currently, as this business is just starting to grow, serving these large enterprises is mostly oriented toward the customer service scenario. But I know we've had great success before. So I'd like to ask you to look ahead - since productivity and production relations are interdependent, like electricity initially being point-based and then transforming entire production processes - might we see, for example, process reengineering at large enterprises? Because currently their sales and service are separate departments. Or might we see other forms of organizations or companies that can better leverage the compound capabilities of AI as you described? Thank you.

BaiRong, Founder, Chairman of the Board, CEO

This question is quite broad and difficult to answer. Frankly, I think: today's BaiRong tries not

to change the client's organizational structure - whatever your process was before, we keep it. Honestly, I visited Silicon Valley in August and came back feeling BaiRong is too conservative internally. Their efficiency is 3 to 5 times ours, so I have a strong sense of urgency.

We're now trying a new model - partnering with consulting firms. Because at this level, BaiRong has an internal system called Silicon Employee Home, which is BaiRong's internal HR system for managing silicon-based employees. It's for internal use only - we never thought about selling it externally. But some consulting firms said, You have this? This is perfect for consulting firms. Because when I explained it to enterprises, they'd say, Show me what an AI digital employee looks like. Seeing BaiRong had this, they ultimately said let's do this together. I said fine.

So we ourselves believe our company's approach is like a caterpillar track. The first wave - programming - we've missed that opportunity, I give up; the second wave we believe is customer service and marketing centers; the third may involve other areas, and we're already starting to enable sales. Some consulting firms have already said your Silicon Employee Home is about managing the collaboration between silicon-based and carbon-based employees, and the ultimate implementation would be transformative.

But that would be the biggest change for any enterprise. I have two examples: a senior executive from a joint-stock bank, and a consulting company person who saw it and said: Actually, the biggest thing that shocked me is your Silicon Employee Home - I've never seen a thousand-person enterprise managed this way. I said compared to Silicon Valley, I'm still not as good as them. This is extremely shocking for enterprise leaders.

So we chose a relatively easy starting point - the contact center - because results are easy to measure, and in most enterprises it's a cost center. BaiRong previously did well because we helped enterprises make money; now touching the cost center, you can see there's already resistance - as seen with the logistics industry client. So sometimes it's three steps forward, two steps back, or even one step forward, two steps back. It's a winding process, so we need to have realistic difficulty expectations. But once we're embedded, replacing us would be quite difficult, because so many links are involved and no one else is willing to do the full-service model like BaiRong. The more links you handle for them, the harder it is to replace you, because you're deeply embedded in their

business processes. Of course, the resistance still exists. But overall, I think today's logistics client has given us a lot of confidence.

BaiRong, IRD

Great, thank you to all the investors for attending. That concludes our discussion today. Thank you everyone.