



Q3 2025 Earnings Prepared Remarks

November 19, 2025

Legal Disclaimer

This document contains forward-looking statements about our future financial performance and other future events and trends, including guidance. These statements are only predictions that are based on what we believe today, and actual results may differ materially. These forward-looking statements are subject to risks, uncertainties, assumptions, and other factors that could affect our financial results and the performance of our business. We will discuss them in detail in our filings with the SEC, including today's earnings press release, and the risk factors and other information contained in the final prospectus relating to our initial public offering. Agora assumes no obligation to update any forward-looking statements we may make on this document.

Zhao Bin, founder, chairman and CEO

Thanks operator, and welcome everyone to our earnings call.

I will first review our operating results from the past quarter.

We're pleased to report our fourth consecutive quarter of GAAP profitability in Q3, supported by double-digit revenue growth and expanding margins. Total revenues in Q3 reached 35.4 million dollars, up 12 percent year over year. Our GAAP net profit for the quarter was \$2.7 million, with a GAAP net margin of 7.8 percent, and we expect our revenue and net profit to continue growing on a quarter-over-quarter basis in Q4.

As you can see, our core real-time engagement PaaS business is rebounding strongly and is on track to deliver its first full-year revenue growth since the pandemic—providing a stable, profitable foundation for us.

At the same time, we're significantly increasing our investment in conversational AI.

Voice-based human-machine interaction is not new, yet most conversational AI solutions today still disappoint users. Why? Because building voice agents that can converse naturally with a human is just hard.

Just a few months ago, Greylock Partners, a leading venture capital firm, published a blog post titled "Voice Agents: Easy to Use, Hard to Build". They nailed the core challenge: behind the simplicity users expect lies immense complexity—system orchestration, real-time audio processing, latency management, and compliance requirements.

Consider the issues of background noise and multiple speakers, just two of the many technical challenges. In real-world settings—like a busy home, office, or car—clean audio is the exception, not the norm. A voice agent must accurately isolate a user's voice from overlapping speech and ambient sounds. Without this, transcription becomes unreliable, intent is misunderstood, and the agent's reasoning falters—undermining the entire interaction.

Furthermore, as Andrey Karpathy has pointed out, there is often a significant gap between a working demo and a production-ready product. Conversational AI is no exception. For instance, in our

discussions with customers and prospects, many have expressed frustration with the reliability and scalability of current solutions, especially when users are distributed across geographies or when concurrent usage is high.

Our investment in conversational AI is specifically aimed at addressing these challenges.

Recently, we launched our Conversational AI Engine 2.0. It integrates over a decade of advanced audio R&D—including AI-powered noise suppression, acoustic echo cancellation, a proprietary audio codec, and adaptation across thousands of device types—to ensure that AI hears and speaks with consistent clarity. In addition, the engine also tackles core interaction challenges: selective attention, turn-taking, interruption handling, emotion detection, and natural conversational flow. In short, we're not just providing the transmission pipelines for voice and video—we're building the behavioral intelligence that powers truly responsive, human-like conversational AI agents.

To help developers build voice agents more easily, we announced our Conversational AI Studio at our recent Convo AI & RTE Conference in late October, which allows developers to create, configure and deploy voice agents through a zero-code interface. Complementing this, our Conversational AI Benchmark & Orchestration Platform allows developers to evaluate, mix and match, and optimize both our proprietary and third-party modules—so they can identify the best-performing combination for their specific use case.

Our open-source TEN Framework—designed for building voice agents—continues to gain traction in the developer community. Recognized for its high-concurrency architecture and deep cross-platform integration, it has been adopted by multiple cloud providers and major enterprises for their agent orchestration platforms.

All these products are backed by our globally distributed, real-time inference cloud. Over the past several months, we've expanded this infrastructure to cover key regions across North America, South America, Europe, and Asia—ensuring consistent latency, reliability, and performance, even under high concurrency and varying network conditions.

Early adoption from customers around the world has been encouraging, and our pipeline of use cases and prospects continues to grow as we head into next year. Our recent Convo AI & RTE

Conference attracted more than 3,000 on-site attendees —a record for us—and made it the largest gathering focused on conversational AI technology globally. Our customers and developers are deploying our conversational AI solutions to build voice agents for outbound marketing, inbound customer service, tutoring, and many other applications. Hardware manufacturers are also integrating our technology into smart toys—enabling voice-powered companionship and learning experiences.

In conclusion, the convergence of advanced AI models and robust, real-time infrastructure is unlocking a new era of possibilities. Backed by proven scalability, deep technical expertise, and a forward-looking product suite, we are well-positioned to empower this next chapter—enabling truly human-like, reliable, and scalable voice agents.

With that, let me turn things over to Jingbo, who will review our financial results.

Jingbo Wang, CFO

Thank you, Tony. Hello everyone.

Let me start by first reviewing financial results for the third quarter of 2025 and then I will discuss outlook for the fourth quarter.

Total revenues for the third quarter reached \$35.4 million, up 12% year-over-year, representing our third consecutive quarter of double-digit organic growth.

If we look at the two business divisions, Agora revenues reached \$18.2 million in Q3, representing 15.9% year-over-year growth and flat quarter-over-quarter. The strong year-over-year growth reflects our successful market penetration and growing adoption in verticals such as live shopping. Shengwang revenues reached RMB122.4 million in Q3, up 8.4% year-over-year and 6.0% sequentially, driven by continued business expansion and adoption in key verticals such as social and entertainment and IoT.

Dollar-Based Net Retention Rate is 108% for Agora and 90% for Shengwang, marking the fourth consecutive quarter of improvement for both businesses.

Gross margin for the third quarter was 66%, slightly decreased 0.7% year-over-year and 0.8% sequentially.

Moving onto expenses. R&D expenses were \$13.8 million in Q3, decreased 52.8% year-over-year. R&D expenses represented 39.1% of total revenues in the quarter, compared to 92.7% in Q3 last year.

Sales and marketing expenses were \$6.5 million in Q3, decreased 5.6% year-over-year. Sales and marketing expenses represented 18.3% of total revenues in the quarter, compared to 21.7% in Q3 last year.

G&A expenses were \$5 million in Q3, decreased 48.4% year-over-year. G&A expenses represented 14.1% of total revenues in the quarter, compared to 30.8% in Q3 last year.

Moving on to the bottom line. We delivered net income of \$2.7 million in Q3, representing a 7.8% net income margin. This result represents a significant improvement from last year and marks our fourth consecutive quarter of GAAP profitability. Based on our current business momentum and visibility into the fourth quarter, we expect

net income to grow sequentially compared to Q3.

Now turning to cash flow. Operating cash flow was \$0.7 million in Q3, compared to negative \$4.6 million last year.

Moving onto balance sheet. We ended Q3 with \$374.3 million in cash, cash equivalents, bank deposits and financial products issued by banks. Net cash outflow in the quarter was mainly due to share repurchase of \$4.8 million.

In the third quarter, we repurchased 5.2 million ordinary shares, or 1.3 million ADSs, representing 1.4% of our outstanding shares at the beginning of the quarter. Since our board approved the share repurchase program in February 2022, we have repurchased \$132.1 million worth of shares through September 30, 2025. The share repurchase program demonstrates our dedication to returning value to our shareholders, balanced with our ability to continue investing in strategic growth opportunities.

Now turning to guidance. For the fourth quarter of 2025, we currently expect total revenues to be between \$37 and 38 million, compared to \$34.5 million in the fourth quarter of 2024, representing year-over-year growth rate of 7.2% to 10.1%. This outlook reflects our current and preliminary views on the market and operational conditions, which are subject to change.

In closing, I would like to express my gratitude to our outstanding teams in Agora and Shengwang. Our sustained double-digit revenue growth and profit expansion are a direct reflection of your hard work and strategic focus. To our shareholders—thank you for your continued trust. We remain focused on executing our roadmap to build a durable, market-leading company at the forefront of AI innovation.

Thank you all for joining today's call. Let's open it up for questions.