

The State of Conversational AI in Support



The State of Conversational AI in Support

How Conversational Voice AI presents an opportunity to 10x support outcomes



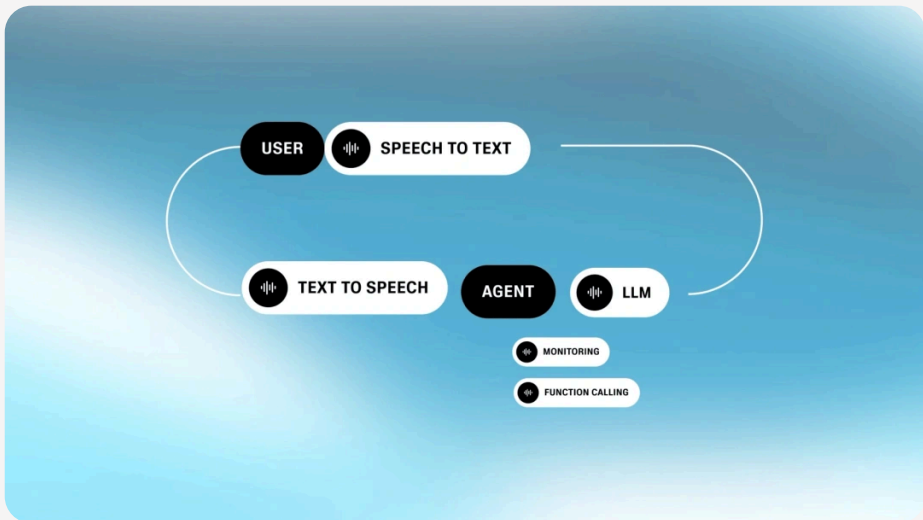
Ben Butler, previously Support Lead at Stripe.



Stripe is synonymous with great support, often going the extra mile to surprise and delight. However, great support is more than just the handwritten notes and 3D-printed totems we see on Twitter. More than a decade's worth of human and software engineering has enabled Stripe to build a customer service offering that supports millions of businesses a year. I joined Stripe in 2015, when the company had just tipped 200 people, and the support team fit around one conference table.

I was one of the first hires in Dublin, an early outpost in Stripe's global support coverage. In my five years there, Stripe re-organized the support team into specializations, introduced external vendors as Tier 1 support, expanded channels, added languages, and, eventually, paid support plans. Having helped design and test these evolving support programs, I saw firsthand the complexities of scaling customer service—complexities that conversational AI could have streamlined.

Conversational AI



Conversational AI has three elements. These three elements work together seamlessly: Speech-to-Text captures the user's intent, Language Models interpret and generate intelligent responses, and Text-to-Speech transforms those responses into natural conversation.

- **Speech to Text:** The AI ears convert a user's spoken natural language into text, with high accuracy.
- **Language Models:** The AI brain processes this text and understands the context to generate intelligent responses.
- **Text to Speech:** The voice transforms text responses into natural-sounding speech.
- **Turn Taking:** A dedicated service monitors for interruptions and ensures the conversation is natural and feels human.

When I joined Stripe, my job was to ingest as much of the product as I could. As part of my training, I would shadow more experienced agents, and learn from a litany of prior support interactions—a similar approach to how an AI agent would learn.

AI is great at pattern recognition, and much of support is a pattern recognition exercise. It starts with a user explaining their issue. The support agent processes that information, tries to match that to their existing knowledge, and presents an answer back to the user. However, support is more than just the recitation of a knowledge base; great customer service requires empathy, creativity, and working together with a user to solve their issue. Unfortunately, most support interactions are found lacking.

The support status quo

Support is an important touchpoint, but it's often painful. Think of the last time you had to call an airline. You're on hold for hours and — when you finally get through to

Conversational AI

Method	Min.	Avg.	Max.
Human (internal)	1.40	3.44	8.19
Human (outsourced)	0.256	0.576	0.96
AI	0.026	0.031	0.036

Source: thinkhdi.com

Conversational AI in action


In the largely analog world of supply chain, Traba is a technological bridge for talent, helping industrial businesses recruit and manage their workers. While Traba provides apps and dashboards, phone communication remains a critical connector for a workforce used to paper timesheets and cork job boards.

Traba looked at what phone calls its operations team spent the most time on and found two dominant areas: fielding support queries or running through scheduling checklists. For CTO Akshay Buddiga, automation was the obvious answer, but timing is also important. “When we identify a process we want to make more efficient, we prefer to embed automation early rather than trying to retrofit solutions later.”

Traba is laser-focused on unit economics and staying ahead of the curve on tech innovation, so moving these calls to conversational AI was appealing for multiple reasons. With speed and cost in mind, the team went looking at solutions. But user experience was also key. As software engineer Joseph Besgen noted: “We wanted it to sound like a conversation, not just listening to a recording.” The ElevenLabs demo sounded so real that, during testing, a Traba employee’s father couldn’t tell whether it was AI or someone from his son’s team.

While there’s an inherent complexity to scheduling thousands of light industry workers, it shouldn’t be as hard to schedule one (1) doctor’s appointment. But have you tried calling your doctor lately? It’s frustrating for patients, and an inefficient use of healthcare staff’s time. In partnership with ElevenLabs, [EliseAI’s](#) assistants are filling the gap in healthcare administration, running everything from appointment scheduling to billing. One hospital now has 86% of its calls handled entirely by AI agents. This results in not only a 66% reduction in cost per call, but also efficiency gains for the administrative staff. EliseAI has also increased access to healthcare for non-English speaking communities.

The future



Some interesting trends emerge when we consider conversational AI growth by sector. EdTech was the first industry we saw embrace conversational AI – as companies could finally provide customized one-on-one tutoring and language learning at an affordable price. Customer support applications quickly followed, as support interactions are particularly well-suited because they are pattern-recognition exercises: the answer exists in a knowledge base, and the AI agent works to match that to the user's question. We also see a rise in verticalized end-to-end applications, particularly in sectors like logistics and healthcare – Traba and EliseAI, for example. Again, these are areas that are repetitive and predictable, so easily handled by an AI agent.

Conversational AI brings support from the atomic realm into the world of bits. Now, support can be a better experience for companies and their customers. Like your cloud provider, your AI support team can scale up and down. Your customers won't ever have to wait on hold, and you save yourself the operational headache of spinning up support centers across the globe.

For the past year, voice agents have been good at conversation and knowledge retrieval, but I think that's just the start. In 2025, I believe AI voice agents will be the standard protocol for inbound meeting scheduling and "product specialist" type support issues. Although they may be limited to knowledge-retrieval calls, this will make a big dent on support volume and free up human teams to do higher-value work.

By 2026, we will have moved from retrieving knowledge to performing actions. The standard AI voice agent will make API calls and connect with third-party apps. They will routinely perform actions like book meetings and issue refunds.

In 2027, I see AI voice agents progressing from support to customer success. I believe things like entire sales deals will be run by an AI agent, possibly on both the buyer and seller side. For many, this seems like the final frontier, but I think it's just the beginning. Deeply contextual and creative tasks, which we once believed were exclusively human, will increasingly become the domain of AI.

Support will be transformed from a cost-center: first neutralizing expenses, and eventually becoming a profit-center for the business. AI voice agents will proactively reach out to customers, reducing churn and increasing customer LTV. Conversational AI provides many of the benefits of human support, but it also comes with perfect recall, tens of languages, and is operational 24/7. The future is coming, and in many ways, it's already here. Imagine instant, empathic, effective support – every single time. Soon, you'll almost want to call an airline, just for the fun of it.