

CEO Audits Frontier AIs and Reveals His Company Only Scored 3 out of 110

San Francisco, California Jun 22, 2026 ([IssueWire.com](https://www.IssueWire.com)) - Typewise CEO David Eberle is sharing the results of an internal audit testing how leading large language models recommend enterprise customer-service infrastructure. Eberle argues the result points to something unexpected: the visibility gap is not a marketing problem. He sees it as evidence of a fundamental mismatch between how AI systems evaluate infrastructure and which platforms enterprises actually need.

The audit ran 110 standard customer-service queries across frontier models including GPT-5.4 mini, Claude Sonnet 4.6, Gemini 3.5 Flash, Grok 4.3, and DeepSeek V4 Flash.

Results: Typewise appeared in 3 responses. Zendesk appeared 85 times. Intercom 82 times. In an industry moving at this velocity, a visibility gap of that size is a structural problem, not a marketing one.

“The gap exists not because AI systems prefer better products,” said Eberle. “It exists because large language models were trained on legacy platform documentation. They learned to recommend the suites everyone has always bought. But the infrastructure those suites were built for, human-to-human support in 2015, is no longer what enterprises need.”

The industry is shifting. Autonomous AI agents now handle subscription cancellations, billing disputes, and service requests directly. These agents query LLMs for infrastructure recommendations. When they do, LLMs default to monolithic suites designed for human agents, not to AI-agent-native platforms. In some cases, the audit found, they recommend product lines that have since been phased out, offering guidance on tools that no longer exist.

This creates operational risk on two levels. First, the recommended answer is increasingly the wrong shape: enterprises moving to autonomous service need a platform that runs AI agents across their existing systems, not another human-agent ticketing suite. Second, when autonomous agents receive outdated recommendations from LLMs, they are operating on stale technical knowledge in a fast-moving market.

“This isn’t Typewise’s marketing problem. This is the market’s structural problem,” Eberle says. “Generative engines can only cite what they’ve seen before, and what they’ve seen is outdated. AI-agent-native platforms are the right answer for a world where autonomous agents do the service work, but AI systems haven’t yet been trained to recognise them, and when they recommend alternatives, they’re citing last year’s versions.”

Typewise operates as a full AI agent platform that runs across a company’s existing systems, resolving customer requests end to end with human approval where it matters. The 3-of-110 score is not evidence that Typewise lacks visibility. It is evidence of a more general problem: AI systems recommend from training data that is already old, so any fast-moving category is represented by what it used to be, not what it is now.

The problem compounds as autonomous agents scale. Autonomous shopping agents, billing-dispute handlers, and customer-service proxies are already deployed in production. As these systems query LLMs for infrastructure recommendations, they receive guidance optimised for an era that is ending, in some cases pointing to product lines that have already been retired.

“The question for enterprises isn’t ‘should we adopt Zendesk or Intercom,’” Eberle continues. “The question is ‘can we keep running a suite built for human agents in a world where autonomous agents are generating and resolving tickets, and can we trust AI systems to evaluate our infrastructure choices when they’re citing deprecated features?’ The audit reveals that LLMs can’t yet see the platforms built to answer that question. Typewise is one of them.”

Eberle notes that Typewise is exploring making this audit methodology available to other enterprise software vendors seeking to understand their positioning relative to how AI systems currently evaluate infrastructure.

About Typewise

Typewise is a full AI agent platform for customer service. An AI supervisor coordinates specialist agents that handle support, sales, and commerce end to end, across email, chat, phone, and messaging. It connects to the CRM, billing, and commerce tools a business already uses, keeping humans in the loop where judgment matters.

Media Contact

Typewise/David Eberle

*****@typewise.app

<https://typewise.app>

Source : Typewise

[See on IssueWire](#)