CASE STUDY

Recall.ai

Turning your most painful problem into an AI platform

Discover:

- Why infrastructure will make or break your AI product
- How your most painful workflow could be your next AI feature
- The sales-led path to real customer learning

Bessemer Venture Partners

Recall.ai: Building the infrastructure behind AI tools and products for video conference recording

Imagine onboarding a new employee at your company using only written information — emails, documents, etc. — and sending them on their way to do the job without a single word of conversation. Even if they were the smartest person in the world, they'd likely be missing some information, and more importantly, context. Excluding your new employee from meetings or one-on-one conversations, for example, would leave a sizable information gap.

This is the insight that led Co-founders David Gu and Amanda Zhu to build Recall.ai, the API infrastructure company that's become the backbone for how thousands of developers capture and process conversation data. Even as AI agents become increasingly sophisticated, they face the same limitation as that hypothetical new employee: without access to the conversations where the majority of company context lives, even the best-designed agent will struggle to navigate tasks with the same ease as a human. For companies like HubSpot, ClickUp, monday.com, and PagerDuty, the Recall.ai team has solved one of the most painful infrastructure problems in building AI products. And they've done it by living through that pain themselves.

Through our case study series, based on conversations with various SaaS and AI leaders, we're exploring the many different routes of those who have successfully navigated building and commercializing leading AI-native solutions — distilling their hard-won insights into actionable frameworks. Building on our conviction in backing Recall.ai, we sat down with CEO David Gu to dive deeper into Recall.ai's journey of turning infrastructure pain into a business by becoming the vendor they wish they had.

Recall.ai's path from pain to platform

Situation	Before Recall.ai, David and Amanda were running a video conference recording product for user researchers and product managers. Despite the early momentum from growing demand, they were spending about 80% of their time on the unglamorous backend work of reliably capturing meeting data.
Challenge	The duo found themselves trapped in a "no-win situation" where infrastructure success was invisible, and a lack of infrastructure was catastrophic, meaning they couldn't win by over- or under-investing.
Solution	When LLMs began advancing, David and Amanda saw an opportunity to solve the problem they faced with their video recording tool. The infrastructure they had already started building for their previous product was extracted into the beta version of their next product.
Result	After finding product-market fit through lived experience and old competitors, Recall.ai evolved into a full API platform now serving some of the biggest names in modern software. Serving over 3,000 companies, Recall.ai has built an infrastructure layer underneath the AI revolution.

Key founder lessons from Recall.ai on building the backend of AI-native products

1. Domain expertise creates compounding advantages.

David and Amanda were among the most knowledgeable people on meeting infrastructure when they started building. Understanding both the developer and their end users meant they could predict cost structures, suggest optimizations, and advise on product strategy — all from lived experience.

2. Sales-led can work for developer products when the goal is learning.

Conventional wisdom says developer tools should be self-serve from day one. Recall.ai did the opposite for the first three years, forcing every customer into a conversation. This built market understanding that couldn't be replicated through analytics or support tickets.

3. Building for developers requires product-minded engineers.

Recall.ai's hiring centered on finding engineers who had both technical depth and a product mindset to talk to customers directly. The skills create an operating model where engineers understand not just how the software works, but how the business works.

4. Infrastructure pricing should align vendor and customer success.

Usage-based pricing eliminates the tension between growth and profitability that plagues many SaaS models. When your costs scale with customer success, pricing becomes both transparent and predictable.

5. The biggest opportunities often lie in the most painful problems.

Recall.ai succeeded not by racing to build the flashiest AI feature, but by owning the unglamorous, complex foundation every AI product needs. The infrastructure layer wins by solving the problem no one wants to touch.

The no-win situation that became a crucial pain point

In 2020, David and Amanda built a video-conference recording tool for product managers, where the team first encountered the problem. During COVID, recording a Zoom call was still novel and socially uncomfortable, especially considering it was many people's first time even participating in a virtual meeting. But as remote work became the norm, the value of captured meetings became clear, and demand for the product grew. The problem? They were spending 80% of their time on infrastructure.

"When you're a product company, you can't win with infrastructure," David explained. "If you do a really good job with the infrastructure, your customers don't notice and instead say, 'Why haven't you built my feature?' But if you don't do a good job with the infrastructure, you lose data, and then your customers are mad at the result."

The technical challenges were immense. "Meetings don't happen randomly — they cluster. Every morning at 9 AM Pacific, millions of people on the West Coast join their daily standups within seconds of each other. It's like Black Friday traffic, but it happens every single day. The infrastructure needs to scale instantly, process data in real-time, and tolerate zero data loss. A single failed recording could mean a lost customer."

David realized they couldn't invest enough in infrastructure to stay ahead, but also couldn't invest less without facing customer churn. They were stuck in what David called a "no-win situation" — one they hated so much they didn't even want to think about the problem anymore because of how painful it was. But that pain eventually turned out to be the signal they needed.

The inspiration for Recall.ai: when infrastructure expertise met an emerging demand in the AI market

In 2022, when GPT-3.5 and other LLMs started getting good at processing unstructured text, two-person startups were building products that would have been impossible just two years prior. "We realized, wow, there's gonna be so much amazing stuff built over the next few years using conversation data and LLMs," David recalls. "But the thing that LLMs don't solve is the painful build and maintenance of infrastructure."

This was their aha moment. As more companies tried to build AI-powered products that needed conversation data, more teams were about to experience the same infrastructure pain that David and Amanda had suffered through. At that point in time, the meeting recording product market was made up of only a handful of companies — and the infrastructure-as-a-service market for conversation capture was nonexistent.

David and Amanda were likely in the top 10 — possibly top five — most knowledgeable people in the world about this problem domain around reliable infrastructure, which revealed to them the opportunity they had to pivot and build a whole new company. Immediately, the duo started sending cold emails to their old competitors about building their infrastructure. The responses validated a lot. Some people ignored them; others were deeply suspicious. But several replied immediately: "Oh my God, I'm so glad you're doing this because this problem really sucks."

Building for builders: Turning infrastructure into a product

David and Amanda didn't wait to build the perfect product — they started selling before they even had a product to sell. Taking the infrastructure they built for their previous company, they began the work of extracting it into a standalone service. After building out the API infrastructure, they launched an alpha program within a month. The early validation shaped everything that followed, and their experience running an end-user product gave them a unique advantage as an infrastructure company: they understood not just their customer (the developer), but their customer's customer (the end user).

"As an infrastructure company, you need to understand the entire chain of causality from your infrastructure to the end user value," David explains. "Oftentimes, when we're in a conversation with a customer, they don't actually know what's good versus what's bad. They're waiting on their customer to tell them."

In early conversations, David and Amanda advised developers on what end users typically want, predicted cost structures, and helped customers optimize their products to reduce transcription costs — all based on lived experience. Over time, Recall.ai's product evolved in three waves:

- **Meeting Bot API was the foundation — the core API for recording video conferences. This solved the immediate problem for companies like David and Amanda's previous product.**
- **Desktop recording SDK launched in late 2025, enabling AI note-taking style recording that could be embedded into any application. Within months, they were seeing hundreds of thousands of installations and deployments.**
- **Coming soon: Mobile Recording SDK, extending the infrastructure to phone calls and in-person meetings.**

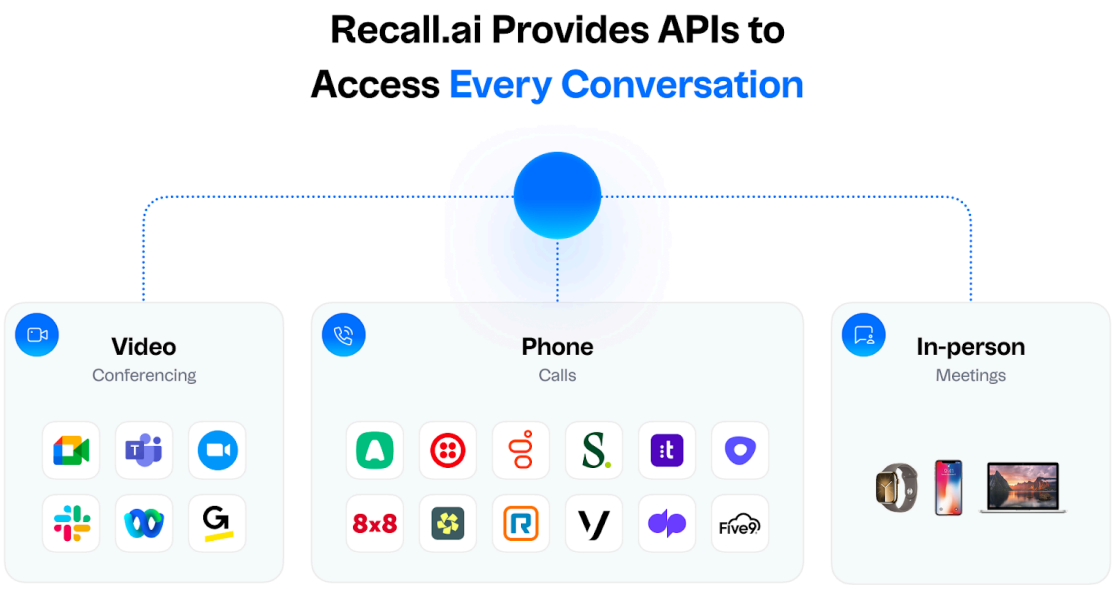

Pricing for alignment: Recall.ai's usage-based philosophy

For Recall.ai, pricing wasn't a complex, strategic decision — it emerged naturally from their infrastructure positioning.

"A usage-based pricing model is very natural for an infrastructure company," David says. "What we say to our customers is if your product becomes very successful and you're now processing a lot of data, Recall.ai is going to cost more. But if you're building internally, testing with your first pilot group of customers, and there's not a lot of volume, then Recall.ai will be less expensive."

This pricing model creates an "alignment of success," where Recall.ai's price aligns with the customer's own revenue. A consumption-based infrastructure model — billing per hour of conversation data captured — also mirrors the kind of pricing developers already understand with Anthropic, AWS, and OpenAI, inherently offering various perks:

- No sticker shock from surprise overages because costs increase only as customer products create more value.
- No artificial friction from seat-based pricing that might limit adoption.
- No need for complex tier structures because usage naturally segments customers by scale.

And this model works across the spectrum — between seven-figure contracts and self-serve pay-as-you-go customers. Elasticity and predictability aren't opposing forces in Recall.ai's model; they're the same thing. Customers can predict costs based on expected volume, and the pricing remains elastic as that volume grows or contracts with their business needs.

Recall.ai's contrarian path to scale: building a sales-led team from scratch

Launching as a two-person meeting bot as a service startup themselves, David and Amanda had to build a team from scratch — a challenging feat, but also advantageous for building an infrastructure company.

"Building the team was a double-edged sword. The positive is that if you find really good engineers on your team, they understand a lot more about the business," David explains. "The hard part is that when we're looking for engineers to hire, we're looking for two traits that are somewhat anti-correlated — a strong backend engineer with a deep understanding of infrastructure, but also product-minded enough to figure out what customers want when designing."

For a developer-focused infrastructure product, Recall.ai's go-to-market strategy should have been obvious: community-first, self-serve from day one, and bottom-up adoption. But David and Amanda did the opposite. For the first three years, every customer had to talk to a salesperson. In the early innings of driving revenue, that salesperson was either David or Amanda.

"The real value isn't the revenue, it's the learning. It's the understanding of what the customers want and why they want it. The challenge with self-serve businesses is that people don't talk to you, and the people who do represent a biased subset — which makes it difficult to develop the zero-to-one understanding of your customer base."

David and Amanda spent three years talking to everyone they worked with. Every conversation followed the same line of questioning:

- What are you trying to build?
- Why are you trying to build this?
- Why do customers want this solution?
- How can Recall.ai help?

Across dozens of verticals — telemedicine, financial services, sales, and customer success — they built a rich understanding of what developers wanted to build. They learned the regulatory constraints in healthcare and why some industries needed phone recording versus video. They also observed that different end users had varying comfort levels with recording technology. In early 2025, Recall.ai launched its self-serve motion, and now thousands of companies sign up weekly.

The learning over revenue trade-off wasn't a sacrifice; it was an investment that's now paying off. Their customer list now reads like a who's who of modern software. But these aren't just customers — they're validation that the world's largest and most widely developed software companies see conversation data as critically important. David is most excited about what this represents: a fundamental shift in how software gets built.

Conversation data as AI's evolving foundation

David's parting insight reframes how founders and CEOs should think about AI agents and automation: "Even if you have a very well-designed agent with a lot of functionality, using the latest and best models, it still needs access to the same context humans do to be able to automate the same tasks."

This is the fundamental challenge facing AI automation today. Even with increasingly sophisticated models, better frameworks for agentic workflows, and more elegant user interfaces, an agent drafting a sales follow-up email still won't know what was discussed in a meeting — and automation falls short of its promise. As AI agents become more prevalent across every industry, data becomes the foundation that determines whether automation works. The companies that figure out how to capture, process, and leverage conversation data will build AI products that feel genuinely helpful rather than limited.

CASE STUDIES

Launching AI products that win

Leaders share how they successfully built AI products. Subscribe to get more insights to help your team win.

Venture insights that matter

bvp.com/subscribe

About Recall.ai:
Recall.ai is the unified API for meeting recording, helping developers capture audio, video, transcripts, and metadata from video conferences and in-person meetings. Founded in 2022 by David Gu and Amanda Zhu, Recall.ai powers conversation data infrastructure for over 3,000 companies. Learn more at Recall.ai.