CASE STUDY

Ada

Powering AI agents with fanatical CX loops

Discover:

- Why your product identity should be rooted in outcomes, not technology
- How to align your pricing model with the evolution of what your customers value
- The value of AI agent resolution depth > resolution rate

Bessemer Venture Partners

Ada: Architecting fanatical CX loops that power AI agents

The early days of building are often inseparable from talking to your intended customers. Every support email is an opportunity to learn, build relationships, or resolve someone's problem. A growing customer base can bring welcomed problems to solve. Still, some complexities compound — support tickets pile up, wait times stretch, and teams shift their focus from answering bespoke questions to categorizing technical challenges and use cases, documenting them, and determining the right communication to solve customer problems.

Ada Co-founder and CEO Mike Murchison felt this transformation acutely at his previous company. "We felt like we were rejecting customer contact," Mike recalls. "Our customers had become these anonymous numbers that we were trying to keep at bay." Avoiding contact doesn't have to become the inevitable cost of scaling your customer base. Instead, he asked: How can businesses encourage and welcome customer contact even at enterprise scale?

That question led to market research: Mike and his co-founder David would become customer service agents themselves, working inside seven different companies' support teams. What they learned across those experiences became Ada — an AI platform now processing 1.5 trillion tokens monthly and powering agents that, for the first time in 2025, consistently outperform human teams.

Through our case study series, Launching AI products that win, we're talking to SaaS and AI leaders to uncover how they successfully built and commercialized leading AI products and features. We sat down with CEO and Co-founder Mike Murchison to unpack how the challenges of scaling customer service at his previous company inspired the idea for an Agentic Customer Experience (ACX) platform, which became Ada, now helping over 350 enterprises deliver exceptional customer service.

Ada's journey from contact center agents to an AI CX platform

Situation	Before building Ada, Mike and David were running a fast-growing company that hit a wall in scaling customer service. The conventional approach of treating support as a cost center to minimize felt fundamentally wrong. They interviewed dozens of CX leaders and learned this was the dominant industry model.
Challenge	The challenge wasn't just operational; it was philosophical. The industry was treating customer service as something to avoid at scale, but Mike believed this couldn't be the future. Incumbent SaaS platforms made it more difficult, not easier, to provide great experiences to customers. But no one had figured out how to maintain personalized, high-quality service while scaling to millions of customers.
Solution	Mike and David joined seven different customer service teams as agents, grinding through tickets in incumbent platforms to understand what made great agents effective. They became top performers manually, then built ML-powered software to replicate those behaviors. The result was Ada's Customer Experience (CX) platform, enabling enterprises to deploy, measure, and continuously improve AI agents across any channel.
Result	In 2025, Ada's customers achieved a breakthrough: their AI agents began consistently outperforming human teams, autonomously resolving 80% of inquiries at higher customer satisfaction rates than human agents, fundamentally changing the economics of customer experience. Ada transformed support from a "cost center" into a driver of customer engagement, product development, and revenue.

Key founder lessons from Ada

1. The best insights come from doing the work yourself, not just observing it.

Mike and David didn't just interview customer experience leaders; they went a step further and asked those leaders if they could join their teams. They worked in the trenches of support desks, experiencing firsthand how the current technology could hinder great service. This immersive experience revealed insights no interview could capture: how the software limitations forced agents to develop workarounds, what made certain agents more productive, and why customers responded differently to various approaches.

2. Product identity should be rooted in outcomes, not technology.

Mike described Ada as "a customer experience company first and an applied AI business second." If another technology emerged that could improve customer experience better than AI, Ada would pursue it. This outcome-first identity protects you from building AI for AI's sake — a trap many enterprise AI companies fall into today. Ground your product identity in the transformation you're enabling for customers, and let the technology be the means rather than the end.

3. The right pricing model emerges from understanding what customers value.

Ada began with outcome-based pricing (charging per resolved conversation), but soon discovered their enterprise customers strongly preferred predictable, conversation-based pricing with annual commitments. Why? Because enterprises adopting AI-native software aren't experimenting anymore — they're going all in and building their businesses around it. They need predictability for budgeting and forecasting, so pricing models must align with customer confidence and procurement realities.

4. Making AI easier to use doesn't reduce demand for AI management (it increases it).

Because people now have more leverage in controlling AI agent behavior, Ada's customers are managing their agents far more actively than before. This is Jevons Paradox in action: when efficiency improvements make AI more accessible, the inference demand increases. If you're building AI products, don't assume ease of use = reduced engagement. Design for power users who will push your platform's limits.

5. Resolution depth, not just resolution rate, determines AI agent value.

In Ada's early days, agents could answer customer questions with a "shallow resolution" that helped users self-serve. But the true unlock came from Ada's "deep resolution," which was doing the work for customers rather than explaining how to. For example, instead of telling a customer how to reset a password, Ada's agents now automatically reset the password and diagnose why the reset was needed. When defining your AI's value proposition, don't stop at making tasks easier; build in-depth workflows that execute automatically.

The customer service problem every business faces at scale

The problem Mike and David identified went beyond the company they were running previously. They talked to VP of Customer Experience leaders and Chief Customer Officers across dozens of companies and heard the same story: the conventional customer service model fundamentally fractures at scale. Most companies often do what every scaling playbook tells them — treat customer service as a cost center to be optimized down. They implement deflection strategies to keep customers from reaching humans, measure efficiency by minimizing handle time, and offshore support or build elaborate self-service portals. The culture shifts from seeking customer contact to avoiding it.

Mike and David saw this transformation firsthand. "We woke up one day, and we said, 'there's no way this will be the dominant customer experience model in the future," Mike recalls. "The average person in North America will waste about 40 days of their life on hold if current trends continue." That's not just inefficient, it's fundamentally antagonistic to what a great customer experience should be like. The founders believed that at some point, businesses would figure out how to scale customer service without sacrificing quality. But achieving that would require completely different technology than what existed.

A light bulb moment: Becoming the customer service agents they wanted to replicate

Mike and David went back to seven of the CX leaders they'd interviewed and asked if they could join their teams — all seven said yes. They instantly gained industry experience and customer empathy in hyperdrive as they worked from a remote office setup in Toronto and started powering through customer service tickets using the incumbent SaaS platforms of the day.

"We just grinded as customer service agents inside the incumbent service products," Mike explains. The experience was revelatory. They weren't just observing the challenges; they were living them. The software they worked with wasn't ideal, being optimized for management metrics like handle time rather than customer satisfaction or agent productivity, but Mike and David still became stellar agents. They worked tirelessly to be among the most productive on their teams, manually, before writing a single line of code. And they saw the impact. What they found was that when they provided more personalized, elevated experiences, customer churn went down and customer loyalty improved.

The convergence of factors to create a successful AI product became clear. Mike and David had:

- Firsthand experience of what made great agents effective
- Access to massive amounts of customer service data
- And machine learning expertise to automate the behaviors that drove great outcomes

Ada's market research was next level because it was in the field with a tight feedback loop. "We invested in trying to figure out how to replicate the behavior that made us great agents manually with software," Mike shared. They took an ML approach because of the amount of data they had, and then focused on making the techniques easy enough for non-technical teams to employ. This is when Ada came to life: not as a theoretical solution, but as software that could replicate the value Mike and David were creating manually with their raw effort.

Their first signal of success? Those seven leaders let them run the first version of Ada inside their teams — and they didn't get fired. "That's when we knew we were onto something, because the value we were creating autonomously with our software mirrored the value we were creating manually." It was the perfect A/B test.

How Ada was built: From shallow to deep resolution

Ada's evolution involved a progression from answering questions to actually solving problems, or what Mike calls moving from "shallow resolution" to "deep resolution." This is how Ada developed their beta product into a full operating model equipped for longer, more complex workflows.

Stage 1: Great informational experiences

The earliest version of agentic Ada was very good at ingesting a company's knowledge base and powering conversational agent experiences across all channels that answered customer questions accurately. If you asked how to reset your password or refund an order, Ada could explain the steps clearly. This was valuable because it helped users self-serve and reduced the load on human agents. But Ada was really just providing information, not resolving any underlying issues.

Stage 2: Deep resolution through automation

Ada soon moved beyond informational responses to actually doing the work for customers. Instead of explaining how to reset a password, Ada's agents identified the customer and automatically sent an email with their new password. Instead of describing how to process a refund, Ada executed the refund. This shift required Ada to move from a conversational interface to a workflow automation platform. They introduced "Playbooks," or AI workflows that can be created almost automatically from existing standard operating procedures. Playbooks allow Ada's agents to take actions across multiple systems simultaneously, executing complex processes that would otherwise require human agents.

Today, Ada's agents autonomously handle complex, multi-system workflows — managing a dozen or more actions per task — saving teams up to 30 minutes per resolution, from a single platform across every customer channel.

Stage 3: Root cause resolution

But Ada didn't stop there; the next frontier was understanding why a customer needed help in the first place. "Resolution depth is something we continue to invest in," Mike shared. "Like for a password reset, what was the issue that led them to request a new password?"
Ada now explores root cause resolution by seeking to answer questions like:

- Why was the customer locked out?
- Did they forget it?
- Or is there a deeper problem with how password resets work with the product?

This in-depth investigation is bringing Ada into the back office in new ways, connecting customer experience teams with product development organizations. It's creating a feedback loop where customer service insights can directly inform product improvements. Building a great product is about finding that initial wedge into the workflow or the feature and then earning the right to expand and grow with complexity and depth.

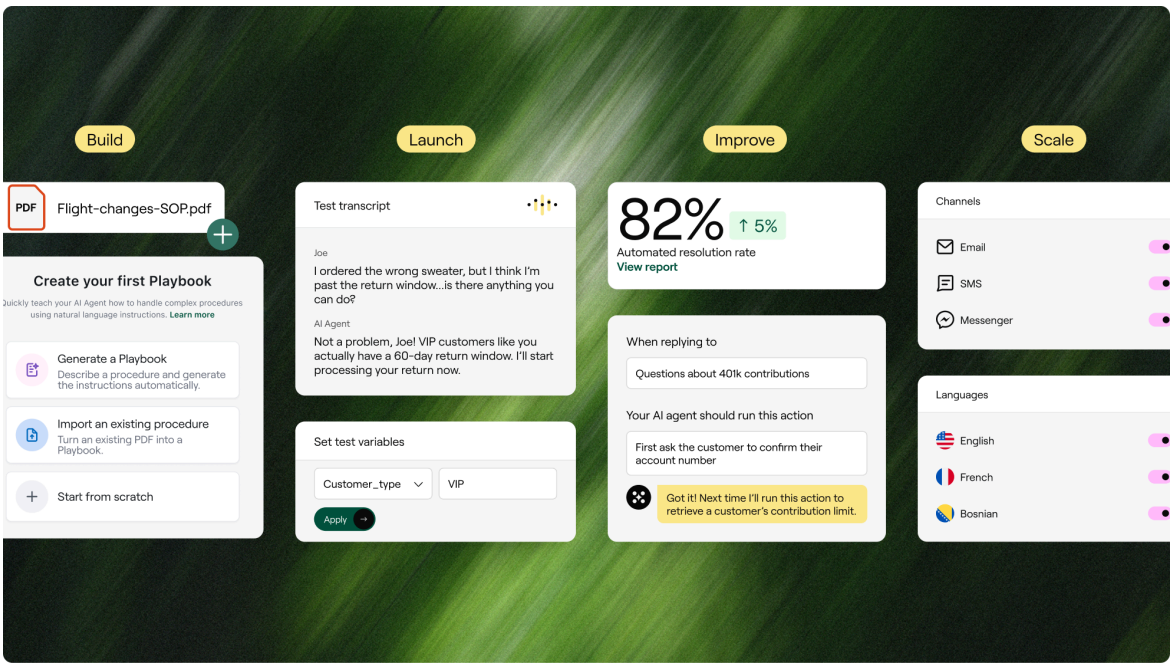

Ada's underlying infrastructure: A unified Reasoning Engine(™)

Powering all of this is what Mike calls Ada's unified "Reasoning Engine," which is essentially a constellation of language models that enables agents to understand customer inquiries, plan resolution paths, and execute them with low latency while adhering to business instructions and brand guidelines.

This constellation is constantly evolving as Ada continuously evaluates and updates which models make up the engine. Increasingly, those models are becoming smaller, more specialized, and fine-tuned on Ada's own data. "If you zoom out on what we do more abstractly, our system is very good at marrying the right intelligence to the complexity of the tasks that a customer's business faces," Mike explained. "We're now living in a time where frontier intelligence is often overkill for a given problem."

The best experiences often come from finding the right combination of models that resolves inquiries fastest with high adherence rates, rather than simply using the largest, most expensive model available. "We've gotten very good at evaluating model performance as it pertains to CX."

Ada's pricing shift from outcomes to predictable conversations

Mike and David's pricing journey revealed to us some important lessons about how enterprise AI monetization is evolving as the technology matures. Initially, Ada went the outcome-based pricing route, charging per resolved conversation or per resolution. This made intuitive sense: AI does the work, so you should pay for the outcomes it delivers rather than the access to the software. For early adopters who were still experimenting with AI capabilities, this model reduced risk.

But as enterprises moved from experimentation to full adoption, Ada discovered that customers preferred purchasing conversations upfront, annually, often with multi-year commitments. They wanted predictability. Why? Because enterprises truly adopting AI-native software aren't experimenting anymore — they're building their businesses around it. Therefore, they need predictable pricing for budgeting and forecasting. They want to align with economic models that support driving more conversation volume, not less, because more engagement with their AI agent is better for business.

"There were a lot of challenges with outcome-based pricing," Mike notes. "I think the fanfare was more reflective of early LLM native software being trialed and folks experimenting with it, and not reflective of a commitment to truly go all in on this software and build your business around it."

Ada's hybrid model today

Ada now operates with a hybrid usage SaaS model where the core unit of value is still a conversation, but those conversations are purchased upfront. Their pricing consists of three components:

1. A platform fee that provides access to Ada's technology.
2. A usage fee reflective of the conversation volume purchased
3. A service or implementation fee reflecting Ada's team integrating the agent deeply into the customer's back office system

This model provides the predictability enterprises want while still aligning Ada's incentives with customer success. Customers who get great results come back to purchase more conversation volume. Mike's perspective cuts through the hype around AI pricing models: "Whether or not the customer is paying directly for a resolution or a conversation, the business case is the same. It's still rooted in an expected resolution rate that we're providing and an incremental improvement in the quality of customer experience."

For founders figuring out their AI pricing strategy, don't get too caught up in the vernacular of pricing units. Focus on creating a strong business case rooted in measurable customer value, then choose the pricing mechanism that best aligns with how your customers actually procure and budget for software.

The go-to-market approach that was primed for the enterprise

Ada's go-to-market strategy reflects a clear conviction about who they're building for and what their customers need to succeed. In Ada's early days, they sold to SMBs and saw great interest in pay-as-you-go pricing and flexible terms. But as the product evolved toward deeper, more complex integrations and workflows, the enterprise became the natural fit. Ada is "very much for the enterprise today," as Mike explains, and this shaped everything about their GTM motion:

- Their pricing is predictable and designed for enterprise procurement
- Their business case is built in partnership with customers, demonstrating long-term ROI
- Their implementation includes services to integrate agents deeply into their customers' back offices

What stands out about Ada's approach is the depth of partnership. This isn't a self-serve product or even a traditional sales-led motion. Ada's team works alongside customers to integrate agents into multiple backend systems. They help translate SOPs into AI workflows and continuously evaluate and optimize agent performance. This hands-on approach stems from Ada's foundational insight: the best customer experiences will come from businesses that learn how to manage their own AI agents, not from outsourcing or handing the keys to forward-deployed engineers (FDEs).

Ada's product identity isn't rooted in chasing every possible use case for LLMs, nor pivoting to the next trend. Their focus and domain expertise on customer experience outcomes allow them to go deeper than generalist AI platforms, building specialized capabilities like their unified reasoning engine and playbook systems that are optimized for customer service workflows.

The future of customer experience: AI agents that exceed human performance

Ada's customers are now achieving:

- Over 80% autonomous resolution rates across all channels (messaging, email, phone)
- Higher customer satisfaction scores than their human counterparts on the same benchmarks
- Over 1.5 trillion tokens processed monthly across Ada's customer base

These aren't occasional success stories. This is becoming the baseline for Ada's most engaged and fully adopted customers. Ada's results prove that the trade-off is breaking down. AI agents can now provide experiences that are both highly personalized and infinitely scalable, handling enterprise-scale volume while maintaining the quality that early-stage companies could once only deliver to their first hundred customers. But reaching this level of performance requires more than just powerful technology. Ada's recent study of 500 enterprise CX leaders revealed that 36% say their teams aren't adequately resourced and skilled to manage, audit, and coach AI agents effectively. Among businesses preparing to transition from human-led to hybrid operations, that figure jumps to 62%. This gap isn't technological; it's operational and organizational.

Their research identified three critical skills gaps preventing businesses from scaling AI in customer experience:

- 28% lack expertise in conversation design and dialogue flow (a core competency for the emerging AI agent manager role)
- 38% report gaps in both system integration and ROI modeling for AI investments
- 80% of businesses surveyed don't yet have a fully adopted governance framework for AI in customer experience, despite governance being essential to safely deploying AI in higher-stakes use cases where errors carry serious consequences

Ada addresses this by helping customers establish dedicated ACX Manager roles, implement measurement infrastructure that separates AI performance from human performance, and build the attribution model that connects AI interactions to downstream business outcomes like retention and customer lifetime value.

With security being a key dimension of the governance framework, Ada's Chief Product and Technology Officer Mike Gozzo explains, "AI agents aren't tools, they're actors. Securing an actor is a fundamentally different problem than securing a tool." For enterprises deploying AI agents at scale, this means treating each agent like an employee: assigning it a managed identity, scoping its access, and auditing its actions.

The North Star of Ada's business — to make the customer experience extraordinary for everyone — is also the driving force for how they're expanding their team, making Ada even better. "The people who join Ada are recruited to build with us," Mike shared. "We're quite obsessive about hiring folks who are driven by self-improvement, share a conviction that the best customer experiences haven't been invented yet, and in many cases, are extremely technical and deeply versed in state-of-the-art machine learning."

For founders building AI products and systems, Ada's story offers an important reminder: the most transformative AI solutions aren't always the ones with the most sophisticated models, but instead are ones that obsess over outcomes customers actually care about and build organizational models that allow customers (both technical and non-technical alike) to own and excel with AI's capability over time.

As the research shows, 92% of businesses expect to increase AI investment in customer experience over the next 12 months, with 21% expecting to operate as primarily AI-led (up from just 9% today). The companies that will succeed in this transition aren't necessarily those making the largest investments, but those making the right ones: building measurement infrastructure, governance models, and workforce capabilities that turn AI ambitions into measurable, repeatable systems. "This technology is evolving at a pace that's difficult to truly internalize, and the future of how businesses communicate with customers is yet to be defined," shared Mike. "We're going to play a major role in defining that."

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 Ada:
About Ada: With more than 550 AI Agents deployed globally, Ada is the trusted leader in agentic customer experience, transforming how enterprises engage with customers. Activated by our ACX Operating Model—a combination of technology, methodology, and expertise—enterprises can easily create and manage high-performing AI agents that deliver personalized interactions across channels.

Since 2016, Ada has powered more than 6.4 billion interactions for global brands like Ancestry, Cebu Pacific, IPSY, monday.com, Pinterest, Square, and Sky, delivering extraordinary experiences at scale. With enterprise-grade security and compliance (SOC 2, GDPR, PCI, HIPAA, AIUC-1), Ada helps organizations reduce cost-to-serve, elevate CSAT, and increase lifetime customer value. Learn more at ada.cx