

CASE STUDY

Papaya Global

Reimagining AI's value by embedding capabilities across an established SaaS platform.

Discover:

- How Papaya built a "family" of AI-powered capabilities
- Why you should operationalize AI for your customers instead of selling an add-on
- Papaya's hard rule for shipping AI

Bessemer Venture Partners

Copyright @ 2026



Papaya Global: When AI stops being a feature and becomes the floor

How Papaya built a family of AI capabilities into its platform	3
Key builder lessons	4
The challenge	5
The problem: global complexity handled at manual scale	
The insight inside the support queue	
The solution	6
Building the AI family: three use cases, one philosophy	
Rethinking ROI when customers won't pay more for AI	8
The result	
Educate, enforce, repeat: the nontypical GTM motion	
After the MVP: Papaya's next chapter	9



Every SaaS company building AI features faces the same tempting assumption: customers will pay more for AI. The logic being that adding intelligence to products means adding a line to the pricing sheet. Papaya Global tried it, and it didn't work. "People who are in the market today expect AI," says Amit Levi, Chief Product Officer at Papaya Global.

For Papaya—the global workforce management platform that helps enterprises hire, manage, and pay employees across 160+ countries—AI became the mechanism for doing something that had always been the company's core promise, just never quite achievable at scale: simplifying the complexity of global employment administration. Companies managing a global workforce navigate countless overlapping compliance regimes, currencies, formats, and time zones. That kind of complexity doesn't compress easily, so Papaya's insight was that AI, applied as infrastructure embedded in their platform instead of a product feature, could finally make that compression more feasible.

Papaya Global's story is part of our case study series, *Launching AI products that win*, where we sit down with SaaS and AI leaders to learn how they successfully commercialized AI. We spoke with CPO Amit to learn more about why Papaya took a distributed, use-case-first approach they call "embedded AI," where they weaved AI capabilities into existing customer workflows across the company instead of launching a standalone AI product. We'll cover the use cases, the key pricing lesson, and the adoption challenges Papaya faced and overcame.

How Papaya built a family of AI capabilities into its platform

Situation	Papaya Global is a global workforce management platform serving large B2B enterprises. Their competitive edge rests on two deeply specialized domains: global employment compliance (regulations, tax, termination rules across 160+ countries) and cross-border payments infrastructure backed by financial licenses.
Challenge	As AI matured, Papaya faced a non-obvious problem: they weren't a data-rich company. Standard AI playbooks (train on huge behavioral datasets, find patterns, optimize) didn't apply as straightforwardly. Meanwhile, customers were fielding an overwhelming volume of country-specific HR and compliance questions that required expert knowledge, making it difficult to answer consistently at scale.
Solution	Rather than building a centralized AI product, Papaya's CPO Amit Levi and his product team took a distributed, use-case-first approach. They mapped Papaya's operational pain points against available AI capabilities, identified the best matches, and embedded AI directly into existing workflows. They explicitly chose not to sell AI as a separate product.
Result	The result is what Amit calls "a family of AI-powered capabilities" woven invisibly into Papaya's platform. Onboarding times dropped from days to hours. Papaya's headcount held flat while the number of workers served scaled significantly, proving AI's value through operational leverage.



Key builder lessons from Papaya Global's story

1. Don't sell AI, operationalize it.

Papaya discovered early that enterprise customers see AI features as a given, not a premium they pay for. The better metric is what AI enables you to do with the same or fewer resources. Measuring revenue per employee and workers served per headcount became Papaya's guiding force, rather than AI-specific ARR.

2. Match AI to your actual data environment, not someone else's playbook.

Most AI success stories come from data-rich companies, yet Papaya wasn't one. Amit's team explicitly mapped their business context (low-volume, high-stakes transactions) against available AI capabilities before building anything. The lesson: start with an honest inventory of where AI can create value between a customer problem and a technical capability.

3. Distributed AI beats centralized AI—embed in the teams that will use it.

Papaya made a deliberate choice to embed AI within individual teams served (legal, operations, compliance, product, marketing) rather than build a centralized AI function. Centralization in a regulated environment creates approval bottlenecks and higher costs. Team-level ownership drives faster adoption and more contextually relevant applications.

4. Every AI release must ship with a feedback loop.

Papaya's hard rule: nothing gets released without a thumbs-up/thumbs-down mechanism. AI that worked in testing can degrade as data patterns shift. Continuous feedback is the only reliable signal of whether the model is still performing in production.

5. Adoption is a change management problem, not a product problem.

People prefer to deal with people. Even when an AI agent can answer a compliance question faster and more accurately than a human can, users will default to reaching out to a person. Papaya addressed this through internal team training (if it's available via self-service, redirect the user and don't answer for them) and external customer education campaigns.



The problem: global complexity handled at manual scale

For enterprises with employees spread across dozens of countries, the problem is twofold:

The first is a knowledge problem.

An HR manager in London needing to hire someone in France sounds simple enough, except the questions stack up: Can this person be brought on as a contractor, or do they need to be a full-time employee? What taxes apply? If the manager wants to pay them 10,000 euros, will the employee actually net 10,000, or are there additional employer contributions on top? And if things don't work out, what does termination look like?

Every country has its own answers to these questions, and those answers change. Papaya had built deep expertise in exactly this kind of global employment compliance, but that expertise lived inside a relatively small group of specialists. Every time a customer had a country-specific question, it generated a support ticket that had to be routed to the right person, answered accurately, and delivered in time to actually be useful. The process was slow and inconsistent.

The second problem is operational.

When onboarding a global employee, the new hire may be in a completely different time zone, possibly even asleep when their first-day paperwork needs to be done. This means they're often hired through a local entity rather than directly, adding layers of complexity. Collecting the right documents for the right jurisdiction means HR managers are manually building onboarding checklists, corresponding back and forth, and doing a lot of waiting. The whole process routinely takes days because the coordination overhead is genuinely difficult.

These were the operational realities of what it means to manage people across borders at scale. But as Papaya grew, the gap between what customers needed and what a human-dependent process could reliably deliver kept widening, so something had to change.

The insight inside the support queue

Papaya Global didn't arrive at AI through a single eureka moment. It started with a spreadsheet exercise.

Amit Levi and his product team sat down and mapped two lists against each other: the operational pain points that were creating the most friction across the business, and the AI capabilities that actually existed (LLMs, anomaly detection, forecasting, and document parsing). Then they looked for the intersections. That framing mattered because Papaya isn't a data-rich company. Unlike e-commerce or gaming businesses, where millions of micro-transactions generate the kind of behavioral signal that AI typically feeds on, Papaya's core business is built on individual, high-stakes hiring events, such as one person, one contract, one-payment-at-a-time transactions.



So instead of asking "what can AI do?" the team asked a more grounded question: where does our business break down, and is there an AI capability that fits the break? The compliance knowledge problem surfaced almost immediately. Papaya's platform serves HR managers who regularly need precise, country-specific information—knowledge that was locked in the heads of a small number of internal experts and buried in a proprietary global employment database. When they looked at that problem next to what a well-trained LLM with structured knowledge access could do, the match was obvious.

The other revelation came the hard way. Early on, the team assumed that AI features, once built, could become a premium revenue line. But when they went to market with AI as an add-on, customers pushed back immediately. The message was unambiguous: AI isn't a feature upgrade they'll pay more for; it's a baseline expectation. That rejection forced a more durable reframe where AI wasn't a product to sell. It was infrastructure for doing more with the same team.

Building the AI family: three use cases, one philosophy

Amit describes Papaya's AI capabilities not as a single product but as a family related by philosophy, varied in maturity, and each one solving a different problem in the business. What ties them together is the approach: start with a real operational pain point, find the AI capability that fits, build an MVP, attach a feedback loop, and iterate until the numbers move.

The compliance knowledge agent

The first member of the family addressed the support ticket problem directly. Papaya's team scanned the volume and content of incoming compliance questions and cross-referenced them against what they called their "global employment hub," which was a proprietary, continuously updated database of employment regulations across more than 160 countries. The result was an AI agent with direct access to that knowledge base, capable of answering questions like "What are pension requirements in the UAE?" or "Can I offer equity to a contractor in India?" in real time. When the agent doesn't have a confident answer, it flags the gap back to Papaya's internal team, who update the knowledge base accordingly. The system only gets smarter with every unanswered question.

Prompt-driven onboarding

The second use case tackled a different kind of friction: the complexity of onboarding workers who are geographically distant, potentially in a different time zone, and being hired through a multi-entity structure. The old process involved manual form creation, emailed document requests, and manual validation, which was a sequence that regularly stretched for days.

Papaya replaced the front end of that process with a natural language prompt. An HR manager can now describe what they need in plain text: the new hire's location, the documentation required, and any company-specific certifications. The system generates a customized onboarding form, sends it directly to the new hire, and uses AI to scan and validate uploaded documents in real time, flagging anything that looks invalid before a human even needs to review. The result, in Amit's framing, is a compression from days to hours. But the more significant change is experiential. The new hire gets a clear, immediate interface. The HR manager gets visibility into a process that previously felt opaque. And the compliance and document validation steps that used to create bottlenecks happen automatically in the background.



The payslip matching agent

The third use case is the most technically iterative of the three, and in some ways the most instructive about what building with AI actually looks like in practice. Papaya distributes payslips across a workforce that spans all languages, where each payslip format is different. Since payslips contain sensitive financial information, matching each one to the correct employee is more than just an efficiency problem; it's a compliance and data integrity problem.

The team built an AI agent to do the matching.

It started by trying to match on employee names, and it failed badly. Names across different languages and transliteration systems don't line up the way a matching algorithm expects—one person might appear as "Emma Zhang" in one system and "E. Zhang" in another, or not at all if the payslip format lists a manager's name alongside the employee's, which is customary in some countries. The team iterated so that if name matching fails, the agent is taught a new rule, including negative rules. For example, in India, the manager's name often appears on the payslip. The instruction: don't treat names as a match candidate.

Amit likens this to "herding cats." You can ask AI to do what you need, and you can explain the logic, but you can't fully control where it goes, and patience isn't optional. The match rate started at 30%. Today, it's approaching 85%. The target is 99%—the threshold at which the process becomes reliable enough to run without meaningful human review in the "red lane" as Amit and the team call it.

Since those first three use cases went into production, Papaya's AI family has grown considerably. The platform now has 14 AI capabilities live, spanning the full lifecycle of global workforce management:

- On the implementation side, they've launched Client Implementation AI that compresses country-specific onboarding from months to weeks
- On payroll, the AI Payroll Cycle Validator, AI Payslip Distribution tool, and Holiday Payment Calendar Planner handle the most error-prone steps in the pay cycle
- On finance, Journal Entry AI, Invoice Validation AI, and SpendGuard bring automation to a layer of the business that has historically required significant manual oversight
- On intelligence, Papaya has launched AI Analytics, Workforce Intelligence, and Workforce 360, giving customers a clearer, faster view into their global workforce than was previously possible

Each capability follows the same philosophy that defined the first three: a real operational pain point, a matched AI capability, an MVP, and a feedback loop. What began as an experiment has now become a methodology.



Rethinking ROI when customers won't pay more for AI

Papaya learned early on that AI in the platform was table stakes for their customers. That realization forced a more fundamental question: if AI couldn't be monetized directly, how should its value be measured?

The team's first instinct was workforce reduction: if AI handles tasks previously done by people, headcount should shrink proportionally. But that metric proved both hard to isolate and, in practice, misleading. AI doesn't eliminate roles in clean, countable increments. "If my team does 15 tasks, but AI can do it, maybe we need fewer team members. And then we can measure it, right? Good AI will reduce two people, and great AI will reduce five people. But we saw it's actually very hard to measure," explained Amit.

What Papaya landed on instead was a set of operational efficiency metrics tied directly to the workflows AI had entered. Across the board, error rates, red-lane versus green-lane workflow ratios, and time-on-screen became the vocabulary of AI value.

The deeper pricing insight is a reframe that goes beyond Papaya. In enterprise SaaS, AI's most durable commercial impact often doesn't show up as a new line item.

It shows up in the metrics that determine whether you renew, expand, and retain, such as customer experience, error rates, and time-to-value. For Papaya, that reframe has a concrete expression: all 14 AI capabilities are embedded within the platform and delivered as part of existing tiers. There are no add-ons and no AI surcharge. The commercial logic is retention-driven: a platform that reduces manual work and compliance risk over time compounds in value with every passing quarter, making it harder to leave and easier to expand. In other words, stickiness is the monetization strategy. Building AI that customers don't notice but can't live without is a stronger competitive position than building AI you just charge extra for.

Educate, enforce, repeat: the nontypical GTM motion

Getting AI into the product was the easier half of the problem for Papaya Global, but getting customers to actually use it was harder. "People prefer people," Amit says plainly. "Even though a human may have no knowledge at all about global regulations and will just read from the knowledge base, people prefer to talk with people. We trust humans more, at least for now." That dynamic shaped their entire go-to-market approach for its AI capabilities. The product and change management work had to happen in parallel:

- Externally, Papaya invested in customer education before expecting adoption. Webinar campaigns walked customers through what the compliance knowledge agent could do and when to use it. Marketing campaigns reinforced the self-service message consistently. The goal wasn't to announce a feature, but to shift a habit.
- Internally, Papaya implemented a policy with real teeth: if a self-service AI capability exists for a customer's question, support and customer success teams were instructed to redirect rather than answer. "You educate your users to go and work with the self-service capability," Amit explains. "It's hard. Not all organizations like it." But the logic held. An internal team that jumps to answer every question inadvertently trains customers to keep asking people instead of AI.



On top of all that was the feedback infrastructure. At Papaya, no AI capability ships without a thumbs-up/thumbs-down mechanism attached to it. This was a condition of release. Continuous feedback is the only reliable signal of whether the model is still performing in production, and it doubles as proof to skeptical users that their input actually shapes what the tool becomes.

This doesn't reflect a typical product launch playbook. But for enterprise AI aimed at non-technical buyers, the GTM motion is less about acquisition and more about adoption. And as it turns out from Papaya's experience, adoption is mostly a change management problem disguised as a product problem.

After the MVP: Papaya's next chapter

Across all three use cases, the pattern was the same: a product manager identified the pain point, worked with engineering to scope the solution, and owned the outcome. There was no centralized AI team and no dedicated innovation lab. AI capability was treated as a competency all engineers were expected to have, not a specialty requiring a separate function.

When the MVP phase was complete and organizational buy-in was secured, Papaya hired a VP of AI. This was their first step toward building a dedicated team to develop proprietary models and embed AI more deeply into the platform's infrastructure. But Amit is clear about the sequencing: the big investment came after the proof, not before. "I don't think doing a big investment without clear value, fast value, is a good idea," he says. "We said, let's start with an MVP. Let's prove to the company it has value. Let's prove to our clients it has value. That was the first year."

Papaya's early AI work was deliberately scoped by starting with proven use cases, building in feedback loops, and proving value to the organization before asking for a larger investment. The results validate the approach. Most tellingly, the number of workers Papaya serves has tripled while internal headcount has held flat. "If we're able to serve more with the current team, that's a very good indication for us."

Enterprise SaaS companies are increasingly being measured against a new kind of competitor: leaner, AI-native startups that appear to do more with less. The temptation is to compete with them directly. Papaya's answer is more grounded: don't try to become a different kind of company. Instead, use AI to do what you already do, at a scale your headcount alone never could. For companies navigating the same pressure, the lesson isn't about any specific use case or model choice; it's about sequencing your launch. The companies that compound on AI success will be the ones who built the internal credibility to keep going.

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 Papaya Global:

Papaya Global is an enterprise workforce management platform that helps companies hire, pay, and manage employees, contractors, and contingent workers in more than 160 countries. The platform combines deep expertise in global employment compliance with a licensed cross-border payments infrastructure, handling everything from payroll and Employer of Record services to contractor management and workforce analytics. Papaya processes more than \$34 billion annually and is trusted by more than 2,000 finance and HR leaders at global enterprises. Learn more at papayaglobal.com.



Venture insights that matter

bvp.com/subscribe