The AI Product-Market Disconnect

How to build AI products people actually need

Executive Summary

Growth-stage SaaS companies are shipping AI products that fail commercially because they've lost connection to real customer problems.

This whitepaper examines the product-market disconnect: the structural gap between what companies build and what customers actually need. This disconnect emerges predictably as companies scale past Series B, driven by translation layers that distort customer insights, solution prescription that replaces problem discovery, and data avoidance that lets executive enthusiasm override adoption reality.

AI doesn't cause this disconnect, but it can accelerate it dangerously. Companies can now move from concept to production 10x faster, which means they can ship 10x more products customers don't need. The result is a portfolio of impressive technology that never becomes revenue.

The companies successfully moving AI from experiment to commercial impact are focused on ensuring they have strategies in place to avoid the product-market disconnect. Only then are they applying AI to validated customer problems with complete go-to-market infrastructure ready.

This whitepaper provides five structural strategies drawn from leaders who've navigated this journey: establishing true product management, creating direct PM-to-customer channels, implementing data-driven reality checks, requiring ruthless problem validation, and building strategic frameworks. You'll find these strategies and tactics aren't necessarily specific to AI product implementation. What is highlighted below are the foundations that make ALL products succeed, but somehow have gotten lost in the drive to release AI-driven features quickly.

The Challenge

Your company is at an inflection point. You've achieved product-market fit. You've raised Series B or C. You're scaling the team. Revenue is growing. You should be accelerating.

Instead, you're noticing symptoms of something broken:

- Your teams are busier than ever, but velocity is slowing. More engineers, longer timelines.
- Sales keeps requesting features customers "need," but those features don't drive adoption once shipped.
- Engineering gets whiplash from constant priority shifts. What was urgent last month is deprioritized this month.
- Customer success is overwhelmed with support requests on features that were supposed to be intuitive.
- Your roadmap is packed, but you're missing revenue targets.

Activity doesn't equal outcomes.
Now add AI to this environment.

Suddenly you're shipping AI products that don't generate revenue. Instead of solving real customer pain points, these products are solving problems executives fell in love with during a demo.

What's actually happening is a product-market disconnect: the gap between what you're building and what customers need.

AI doesn't cause this disconnect. But it exposes it brutally and accelerates it disastrously. You can now build the wrong things 10x faster, creating a portfolio of impressive technology that never generates commercial impact.

Why This Matters Now

The AI wave isn't waiting for you to fix your product strategy first.

Every one of your competitors is launching AI products. Your board is asking about your AI roadmap. Sales is hearing "what's your AI strategy?" in every deal. The pressure to ship something AI-powered is immense.

This pressure is dangerous when combined with the product-market disconnect. Here's why: AI amplifies velocity in whatever direction it’s pointed. If you're disconnected from real customer problems, AI tools don't fix that. They simply let you build and ship the wrong things much faster.

The cost of failed AI products is escalating. Each product that fails to generate revenue doesn't just waste engineering time, it erodes trust. Your board loses confidence in your product strategy. Your sales team stops pitching AI capabilities. Your customers become skeptical of your "innovation." Your best engineers get frustrated building products nobody uses.

The window is closing. Companies that solve the disconnect problem first will pull ahead decisively. They'll ship fewer products, but those products will generate revenue. They'll have pricing models, customer education, and support infrastructure ready. They'll capture market share while competitors are still wondering why their AI features aren't selling.

You can't AI your way out of a strategy problem. Companies that haven't solved this disconnect first will simply fail faster in an AI-driven world.

Analysis: How AI Makes It Harder

AI tools promise to help you build features faster. And they deliver. You CAN ship products 10x faster with AI assistance. But speed without strategy is just expensive confusion.

The Pressure to Build Before Validating

"The industry has fallen into 'solution looking for a problem' where teams are building AI features because the technology exists rather than identifying genuine user needs first." This observation from a CTO with 30 years of experience captures the core issue.

Companies ship AI products because of external pressure, not because they've identified customer problems that *AI uniquely solves*:

- **Board expectations:** Every investor wants to hear your "AI strategy"
- **Competitive FOMO:** Competitors are announcing AI features, so you feel compelled to match
- **Sales urgency:** Your team is hearing "what's your AI story?" in every enterprise deal
- **Engineering enthusiasm:** Your best developers are excited about the technology and want to experiment

This creates what one executive called "building for VCs, not customers." Companies are scaling AI initiatives before establishing genuine product-market fit, chasing investor narratives instead of customer value.

The Demo Success Trap

AI products often get greenlit based on internal demos that mask fundamental problems. Demos impress. The technology works. Executives love the presentation. Customers express "interest." You've proven it's technically feasible.

The problem is that "interest" doesn't equal revenue. As one seasoned CTO emphasized: "The only validation that matters is customers paying money."

The gap between demo success and production impact reveals what's missing:

- **No validated customer problem:** You built something cool, but not something customers need enough to change their workflows and pay for it.
- **No value metrics:** You can't measure impact because you never defined what success looks like.
- **No pricing strategy:** How do you charge for this? Per-seat doesn't work. Usage-based? Flat add-on? You haven't figured it out.
- **No sales enablement:** Your sales team doesn't know how to position it, demo it, or overcome objections about it.
- **No support infrastructure:** Customer success doesn't know how to onboard users or troubleshoot issues.
- **No change management plan:** Customers need to retrain teams and modify processes, but you haven't helped them with that transition.

AI Washing

"Everyone pitching to VCs needs 'AI' in their deck," noted an industry executive who's seen billions in venture funding create tremendous noise. This makes it difficult to distinguish genuine value from what the industry calls "AI washing" (wrapping existing functionality in AI branding without adding real value).

Companies are launching AI features that are essentially "ChatGPT with a sidebar" rather than solving actual customer problems in differentiated ways. Chatbots that seemed impressive in demos become tedious in daily use when users realize they're just typing requests instead of clicking buttons.

The challenge is evaluating whether solutions solve real-world problems or simply represent ideas searching for problems. Many companies have built "cool" AI features that hit walls in adoption because they don't fit actual user needs or workflows.

The Missing Middle Problem

Research from PWC, Deloitte, and others reveals a consistent pattern in AI adoption: executives understand its importance, frontline workers embrace it, but middle management doesn't get it.

One IT director identified why: AI solutions are often presented as technology adoption initiatives rather than as solutions to specific customer problems. When the value proposition isn't clear to the management layer responsible for implementation, they don't understand WHY they're building it. Without that clarity, products ship without the organizational alignment needed to make them successful.

This explains why you can have:

- Executives excited about AI strategy

- Engineers building impressive prototypes
- Customers expressing interest in capabilities

...yet products fail commercially, because the people responsible for operationalizing the solution don't understand what customer problem it solves. The features make it to production, but without the GTM infrastructure, customer education, and support systems needed to generate revenue.

What Great Leaders Do Differently

The leaders who successfully ship AI products that generate revenue don't just "talk to customers more." They've built structural solutions to structural problems. These are the same foundational solutions we've always used, and that smart leaders stick with regardless of whether the product is AI-driven.

The strategies offered below are inspired from the countless conversations I've had with intelligent leaders in AI-driven SaaS organizations. They solve both the chronic disease (product-market disconnect) and its acute AI symptom (products that fail commercially). Here are some stories that provide examples of strategies that work:

Strategy 1: Establish True Product Management

The problem: PMs weren't engaging with customers. Engineering didn't understand the product. The gap between customer reality and product decisions was killing them, and AI products were dying in that gap.

What they built:
This team established a proper product management function positioned between customer-facing teams and engineering, with clear responsibility for:

- Educating engineering about customer context and workflows
- Translating customer observations into product requirements
- Bridging customer reality to technical constraints
- Managing customer insight flow as a systematic process, not ad-hoc conversations

Why it works:
When AI products fail commercially, it's often because no one translated "customers think this demo is cool" into "here's the actual workflow change required, the pricing model that works, and the support infrastructure needed to scale this."

Proper product management makes customer understanding a core function, not everyone's side job. It creates accountability for requirements quality and builds institutional knowledge

about how customers actually use products, which is the foundation missing from most AI product launches.

Warning signs you need this:

- Requirements are accepted without customer validation.
- Engineers are surprised by how customers use features post-launch.
- Customer success teams are constantly working around product limitations.
- AI products are shipping without clear paths to revenue.

Strategy 2: Direct PM-to-Customer Channels

The problem: Multiple translation layers distorted customer problems into solution prescriptions. By the time information reached product decisions, the original problem was unrecognizable.

What they built:
Made direct customer contact mandatory for PMs. They didn't just host scheduled interviews, they made real observations of actual usage. PMs became responsible for understanding:

- What customers were trying to accomplish (not what they requested)
- Where they got stuck (friction points)
- What workarounds they created (signals of missing functionality)
- What they said they wanted vs. what they actually used

Why it works:
You can't build the GTM infrastructure for an AI product (pricing, sales positioning, customer education, support processes) if you don't understand the actual customer workflow it needs to fit into. Direct observation reveals the gap between "this demo is interesting" and "I'd change my workflow and pay for this."

How to implement:
Start with 2-3 PMs requiring 5 customer interactions monthly:

- 50% observation (watch them work)
- 30% problem exploration (what's frustrating?)
- 20% solution validation (would this help?)

Create feedback loops: PMs share insights weekly. Those insights inform roadmap reviews. PMs track which insights led to successful features.

Warning signs you need this:

- PMs can't explain why features matter in customer terms.
- Roadmap debates are focused on opinions, not customer evidence.
- You have a high percentage of shipped features with low adoption rates.
- Sales is constantly bypassing Product to request features.

Strategy 3: Data-Driven Reality Checks

The problem: Success metrics, like executive enthusiasm and demo interest, masked the production reality which resulted in low customer conversion, usability issues, and ineffecient customer support.

What they built:
Established "data as source of truth" culture with monthly reviews: "What did we believe that turned out to be wrong?"

When enthusiasm conflicted with adoption data:

1. Validate data accuracy (is measurement correct?)
2. Explain the discrepancy (why does reality differ from expectations?)
3. Adjust strategy based on reality, not wishes

Why it works:
Products fail commercially because companies mistake demo interest for production viability. Data-driven reality checks force the question: "Are customers actually using this enough to justify the GTM investment required to scale it?"

How to implement:
Establish baseline metrics *before* products launch:

- Feature adoption rates (% of customers using X within 30 days)
- Usage depth (how extensively do they use it?)
- Customer-reported value (surveys, NPS tied to specific features)
- Support ticket volume per feature

Create review rituals:

- Monthly: Actual vs. predicted adoption
- Quarterly: "What surprised us this quarter?"
- Post-launch: 30-60-90 day feature retrospectives

Warning signs you need this:

- Decisions are justified by "gut feel" without data.
- Data is presented but consistently dismissed.
- The company is surprised by product performance (good or bad).
- The company can't predict which products will succeed commercially.

Strategy 4: Ruthless Problem Validation Before Solution Building

The problem: Teams built AI products because the technology was cool, not because they'd validated customer problems. Months later, shipped products had no path to revenue.

What they built:
Stopped accepting solution-based requests ("we need an AI chatbot") and required problem articulation: "What customer outcome are you trying to enable?"

Validation checklist before any AI initiative entered the roadmap:

- Can we articulate the customer problem in their words?
- How many customers have this problem?
- How are they solving it today (workarounds)?
- What would measurably improve if we solved it?
- Would they pay more / churn less / buy faster if we solved it?
- Why is AI better than their current workaround?

Why it works:
As one CTO emphasized: "Focus on getting people on base rather than swinging for home runs. Interest and enthusiasm don't equal revenue. The only validation that matters is customers paying money."

This validation process surfaces what's missing from most AI products: validated willingness to pay, clear value metrics, and evidence the solution is better than current workarounds. This is the foundation for building GTM infrastructure.

Warning signs you need this:

- AI features are described by technology, not outcome ("AI-powered dashboard").
- Roadmap debates are focused on "cool" vs. "boring."
- Your team can't articulate revenue impact of AI initiatives.
- You have a high percentage of products with disappointing adoption.

Strategy 5: Strategic "No" Frameworks

The problem: Without frameworks for prioritization, every AI initiative seemed equally urgent. Teams built everything, exhausted engineering capacity, and couldn't focus on the GTM work needed to make any single product successful. They needed to learn when and how to say "No" to initiatives that didn't move the needle.

What they built:
Explicit prioritization frameworks that could withstand pressure:

- **Investment thesis:** Every feature must map to one of 3-5 strategic themes
- **Capacity allocation:** Fixed % of capacity for categories (40% core functionality, 30% competitive parity, 20% innovation, 10% debt)
- **Opportunity cost framing:** "Saying yes to X means saying NO to Y. Which matters more?"

Why it works:

Most companies have 5-10 AI initiatives but lack the capacity to properly launch any of them. Strategic "No" frameworks force the question: "Which product do we take to market this quarter with full GTM support?" rather than "Let's ship three more features."

This creates the focus required to build complete GTM infrastructure (pricing, sales enablement, customer education, support processes) for the products that matter most.

How to implement:
Define 3-5 strategic themes, then require every AI initiative to map to one theme. If it doesn't map, it's a strong presumption of "No." If it does, it competes with other initiatives in that theme.

Set capacity budgets per theme to prevent AI innovation from consuming 80% of engineering while core product suffers.

Warning signs you need this:

- Your roadmap contains 10+ AI initiatives at various stages.
- Your team can't articulate what you're NOT building.
- Teams complain about constant context switching.
- There are multiple products shipping, none generating expected revenue.

The Pattern

Notice what these strategies have in common? They don't solve the "AI problem." They solve the product-market disconnect problem that makes AI products fail commercially.

Companies that build these structural foundations first can ship AI products that generate revenue because they have:

- Validated customer problems (not just demo interest)
- Clear value metrics (not just technical feasibility)
- Pricing strategy (not just "we'll figure it out later")
- Sales enablement (not just engineer enthusiasm)
- Support infrastructure (not just impressive technology)

The companies stuck with failed AI products are trying to AI their way out of a strategy problem. The companies winning are *solving strategy first*, then applying AI to validated customer problems with complete GTM infrastructure ready.

Conclusion

If you've recognized your company in these patterns, you're not alone. The product-market disconnect affects nearly every growth-stage SaaS company that scales past Series B. The

difference between companies with failed AI products and those shipping AI features that generate revenue is the presence of strategic product leadership that bridges the gap.

The good news is these structural problems have structural solutions. The strategies outlined in this whitepaper are battle-tested approaches from leaders who've successfully navigated this journey. But implementing them requires more than reading about best practices. It requires someone who can:

- Diagnose where your specific product-market disconnect is occurring.
- Build the organizational structures that reconnect product decisions to customer reality.
- Establish the GTM foundations that turn promising technology into revenue-generating products.
- Create accountability systems that prevent the disconnect from recurring as you scale.

This is the work of strategic product leadership, and it's exactly why fractional Chief Product Officers exist.

I work with growth-stage SaaS companies working toward or at Series B and above, ($10M-$50M ARR) and who are navigating the exact challenges described in this whitepaper. My clients typically come to me when they realize their AI products aren't generating revenue, their engineering velocity is slowing despite adding headcount, or their roadmap is packed but revenue targets are being missed.

My approach is straightforward: a diagnostic phase to identify your specific disconnect patterns, followed by hands-on advisory to guide implementation of the structural solutions that bridge the gap. I don't write strategy documents that sit on shelves. I embed with your team to help build the product management discipline, customer connection systems, and strategic frameworks that turn AI technology products into commercial success.

If your company is struggling to turn AI investments into revenue, the solution isn't to build faster. It's to *solve the strategic foundation first*.

Let's talk about what that looks like for your company.

About the Author

Natalie Cervantes is a Fractional Chief Product Officer specializing in growth-stage SaaS companies navigating hypergrowth challenges. With 20 years of software product and engineering leadership experience across startups, Fortune 100 companies like Microsoft and Cruise, and consulting for brands including Amazon, Trimble, and Coca-Cola, she helps companies transform from feature-driven development to strategic, outcome-focused product discipline.

Through NC Consulting, Natalie works with Series B+ SaaS companies ($10M-$50M ARR) to solve the product-market disconnect that prevents AI products from generating commercial impact.

Natalie Cervantes
Founder & Fractional Executive Consultant
NC Consulting
🌐 nataliecervantes.com
www.linkedin.com/in/nataliecervantes
📧 natalie@nataliecervantes.com