

INTRODUCTION

Your Buyer Just Hired an Agent

Sometime in late 2024, a product marketer we know—let's call her Sarah—noticed something strange in her pipeline data. Sarah ran product marketing for a mid-market analytics platform, the kind of company that sells to data teams at Fortune 2000 firms. She was good at her job. She'd built the positioning, trained the sales team, published the battlecards, ran the launches. Her funnel metrics had been steady for two years. And then, over the course of about six weeks, three things happened simultaneously.

First, her website traffic went up—but her form fills went down. More people were visiting the product pages, but fewer were downloading the white papers and requesting demos. Second, her sales team started reporting that prospects were showing up to discovery calls already knowing things they shouldn't have known—competitive pricing, technical limitations, integration details that lived in documentation Sarah's team hadn't even promoted. Third, and most unsettling, two enterprise deals closed in under three weeks. That had never happened before. The typical sales cycle for Sarah's product was eleven weeks.

Sarah's first instinct was that something was broken in the analytics. Maybe a tracking pixel had died. Maybe the CRM had hiccupped. But the data was clean. What had actually changed was something far more interesting: her buyers had started using AI agents to do the research, the comparison shopping, and the vendor pre-qualification that used to take human buying committees weeks of meetings, spreadsheets, and hallway conversations to accomplish.

The agents weren't doing anything magical. They were reading her product documentation—all of it, including the stuff buried three clicks deep that no human prospect ever found. They were pulling her pricing from G2 and Gartner Peer Insights and running it against competitors. They were synthesizing analyst reports and customer reviews into concise briefs for their human principals. They were, in effect, doing the job of a procurement analyst, a technical evaluator, and a competitive intelligence researcher all at once, and they were doing it in hours instead of weeks.

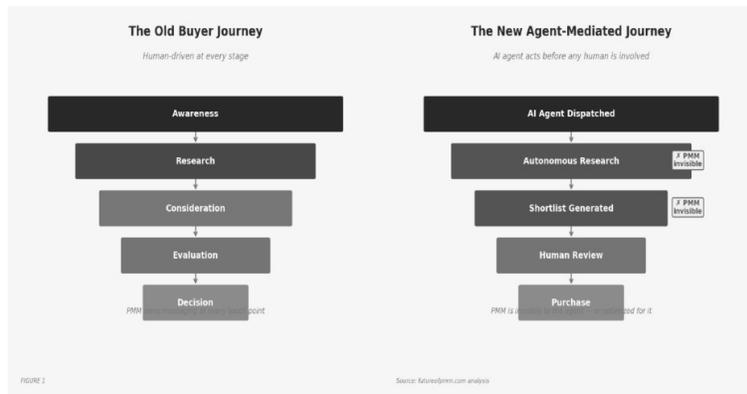


Figure 1: The Old Buyer Journey vs. The New Agent-Mediated Journey. PMMs are invisible during the stages that now matter most.

Figure 1: The Old Buyer Journey vs. The New Agent-Mediated Journey

Sarah's content hadn't changed. Her positioning hadn't changed. Her product hadn't changed. But her buyer had fundamentally changed—or rather, her buyer had hired an intermediary that didn't care about brand affinity, didn't respond to emotional storytelling, and couldn't be schmoozed over a steak dinner. And that intermediary was making decisions—or at least dramatically shaping the decisions—that determined whether Sarah's product made it to the shortlist.

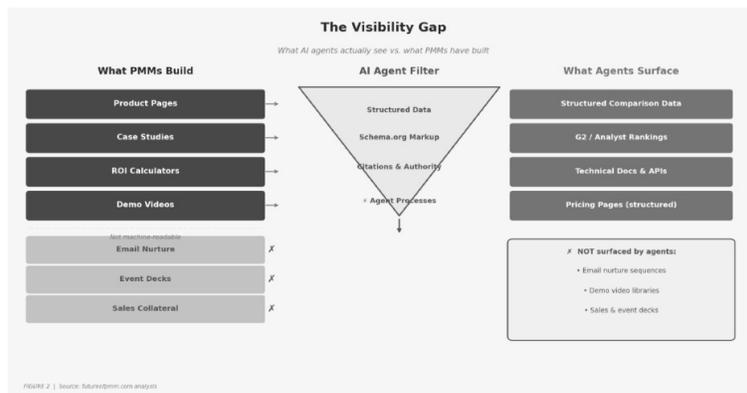


Figure 2: The Visibility Gap. Most PMM content investment is invisible to AI purchasing agents. The assets that agents weight most heavily have historically been treated as operational, not strategic.

Figure 2: The Visibility Gap

• • •

This book is about what happens next.

If you're a product marketer—whether you carry the title of PMM, Director of Product Marketing, or VP of Go-to-Market—you already feel the

ground shifting. You've probably been told that AI will "supercharge" your workflow. You've been handed a ChatGPT license and told to figure it out. Maybe your company has an "AI task force" that meets biweekly and produces PowerPoint decks about the art of the possible. None of that is what this book is about.

This book is about the structural transformation of product marketing as a discipline. Not "how to use AI tools"—though we'll cover the tools, because you need to know what's actually good and what's vaporware. The real story is that the economics, the workflows, and the buyer dynamics that have defined product marketing for the last twenty years are all changing at once. The Pragmatic Institute has a framework with thirty-seven discrete PMM activities. Within three years, the majority of those activities will be performed primarily by AI agents—or they'll be so heavily augmented by agents that the human contribution shrinks to judgment, strategy, and taste. The PMM who could do all thirty-seven at a B-minus level was perfectly employable in 2023. By 2027, that person has been replaced by a person with an agent stack who does eight of those activities at an A-plus level and lets the machines handle the rest.

That's the "10x yourself" thesis of this book, and we want to be precise about what it means. It does not mean "work ten times faster." It means that the output gap between a PMM who has internalized the agentic shift and one who hasn't is going to widen to a factor of ten—in quality, in speed, in strategic impact. The PMM who builds an always-on competitive intelligence system that monitors fifteen competitors in real time is not just "faster" than the PMM who updates battlecards quarterly. They are operating in a fundamentally different category.

willing to state plainly: approximately six percent of B2B marketing organizations are operating in a way that is meaningfully adapted to the agentic buyer environment. The rest are somewhere on a spectrum that runs from 'unaware the transition is happening' to 'has purchased AI tools but hasn't changed the underlying strategy.' The six percent who are thriving share a specific set of characteristics — they have made all three of the moves described in Section 4 — and they are pulling away from the field in ways that will be very difficult to close once the gap compounds for another twelve months.

"Approximately 6% of B2B marketing organizations are meaningfully adapted to the agentic buyer environment. The rest are optimizing for a buyer who no longer exists."



Figure 3: The PMM Readiness Spectrum

• • •

A Word About Who We Are

We've spent our careers at the intersection of data, marketing, and

enterprise software—and between us, we've seen every era of this evolution from the inside. Chris O'Hara, my co-author, was there for the DMP era—he ran global data strategy at Krux, the data management platform that Salesforce acquired in 2016. He was there for the CDP era—he co-authored the book on customer data platforms with Martin Kihn while they were both at Salesforce. And we've both been watching the agentic era take shape from inside SAP, where we lead the product marketing organization for SAP Business Data Cloud. Chris has written three books about the evolution of data-driven marketing, and every one of them was really about the same thing: what happens to business when a new data architecture changes the rules of engagement.

This time, though, the architecture isn't just changing how companies manage data—it's changing who makes decisions, how buying happens, and what it even means to "go to market." That's not a DMP-to-CDP kind of upgrade. That's a wholesale rewiring of the discipline.

As CMO of SAP Business Data Cloud, where we run one of the largest product marketing organizations in enterprise software, we bring the view from the person who has to build the team, approve the tools budget, justify the headcount, and deliver the pipeline numbers in a world where the playbook we learned is being rewritten in real time. Chris brings the practitioner's lens—he's been in the trenches across every era of data-driven marketing, from Krux to Salesforce to SAP. Together, we've tried to give you both angles: the executive's operating reality and the practitioner's playbook. The "CMO Perspective" sections throughout each chapter are where those two views converge—what each topic looks like from the leadership seat, along with the key takeaways we think matter most.

They have made three specific moves — what this chapter calls the PMM Response Framework — that, taken together, mean they are visible to the agent layer, competitive in the data layer, and irreplaceable in the narrative layer. The moves build on each other, but they can be sequenced. The mistake most organizations make is trying to start with Move 3.

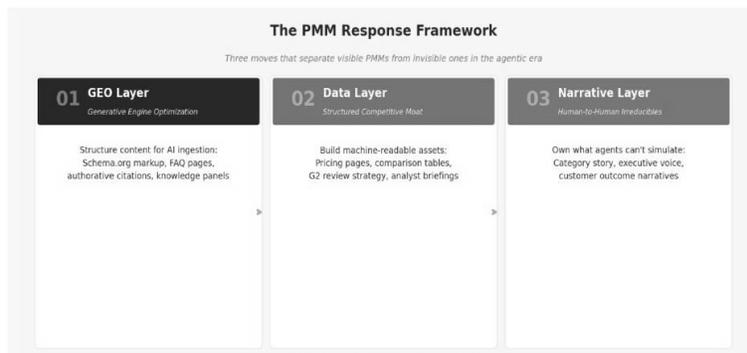


Figure 4: The PMM Response Framework



The Structure of This Book

Part I gives you the macro context. If you've followed the evolution of enterprise data from warehouses to lakes to whatever we're calling the current thing, you'll recognize the pattern—a new data architecture emerges, it reshapes the go-to-market motion, and the marketers who ride the wave early build careers while the ones who wait get commoditized. Chapter 1 tells that story quickly. Chapter 2 takes the Pragmatic Institute's PMM framework—the closest thing the discipline has to a canonical job description—and subjects it to an honest audit. Which of those thirty-seven activities are about to be automated? Which ones become radically more efficient? And which ones become more valuable precisely because they're the activities that machines can't do?

Part II is the meat of the book. Eight chapters, each covering a major cluster of PMM activities: competitive intelligence, positioning and messaging, sales enablement, content and thought leadership, demos and product storytelling, pricing and packaging, launch management, and analyst relations. Each chapter follows the same basic pattern: here's how the activity works today, here's what changes when agents enter the picture, and here's what the best PMMs are already doing about it. These chapters are practical. Named tools, real workflows, specific guidance you can take into your Monday morning. They're also grounded in stories—real people at real companies solving real problems, because that's how we've always approached things and because abstractions are where good advice goes to die.

Part III zooms out to the organizational level. How do you build a PMM tech stack that's actually coherent? What skills should you be developing—and which ones should you stop worrying about? How should a CMO structure a product marketing team in 2027? And what does the PMM career look like in five years—is it bigger or smaller, more strategic or more operational, more specialized or more generalist?

Every chapter in Part II also includes a "Pragmatic Remix"—a mapping back to the specific Pragmatic Institute activities that the chapter addresses, showing concretely which activities are automated, augmented, or elevated. If you're a Pragmatic-certified PMM, you'll see your training reflected back at you through an agentic lens. If you're not, you'll learn the framework as you go.

• • •

One More Thing Before We Start

This book was written with AI.

Not in the way you might think. We didn't hand Claude a prompt and publish what came back. But we used AI tools extensively in the research, the synthesis, the competitive analysis, and the early drafting of every

chapter. We did this deliberately, because the thesis of the book demands it. If we're going to tell you that a product marketer can 10x themselves with the right agent stack, we should probably demonstrate what that looks like. The stories are ours. The judgment is ours. The bad jokes are definitely ours. But the process of creating this book is itself an example of the agentic workflow we're describing—and we think being transparent about that is more honest than pretending we wrote every word in a cabin in Vermont with nothing but a legal pad and a pot of coffee.

Alright. Let's get into it.

THE CMO PERSPECTIVE

The conversation about AI inside marketing departments right now is almost entirely wrong. It's focused on efficiency—how do we produce the same outputs with fewer people, or the same people in less time. That's a CFO question, not a CMO question. The CMO question—the one we wrote this book to answer—is how do we produce fundamentally different outputs? How do we shift from quarterly competitive reviews to continuous intelligence? How do we move from static messaging frameworks to dynamic positioning that adapts to what agents are surfacing in real time?

What we've seen firsthand, running a product marketing organization of fifty-plus people at SAP, is that the PMMs who are thriving aren't the ones who adopted AI tools first. They're the ones who rethought their workflows first and then found the tools that matched. That distinction matters, and it's one of the organizing principles of everything that follows. Throughout this book, we'll toggle between the practitioner's view and the executive view—what a 10x PMM looks like from the inside and what a CMO is looking for when they see it happening.

KEY TAKEAWAYS

- The agentic shift is a structural transformation of the PMM discipline, not just a toolkit upgrade.
- Buyer behavior has already changed—AI agents are compressing research cycles from weeks to hours.
- The "10x" thesis isn't about speed. It's about the widening output gap between PMMs who internalize the shift and those who don't.
- This book covers every core PMM activity through the lens of what's automated, what's augmented, and what becomes more irreducibly human.

The Three Eras of Enterprise Data

There is a company in Stuttgart—we can't name them, but you'd recognize the logo—that has been running SAP for thirty-one years. Their SAP Business Warehouse holds over two decades of transactional history: every purchase order, every invoice, every production run, every quality inspection, every customer interaction, organized into cubes and InfoProviders that a generation of consultants built and a generation of analysts learned to query. When we sat down with their CIO in late 2024, she said something that has stayed with us ever since: "We have more data than any company in our industry. And we have less insight than a startup with a Snowflake account and two data engineers."

That sentence captures the paradox that has defined enterprise data management for the last decade, and it's the reason this chapter exists. The history of how companies collect, organize, and activate data is not just a technology story—it's the story of what becomes possible when the architecture changes. Every time the data foundation shifts, the go-to-market motion shifts with it. The companies that recognize the shift early build advantages that compound for years. The ones that wait get commoditized—not because they lack talent, but because they're running a modern playbook on yesterday's infrastructure.

We think of this history in three eras, defined not by the marketing tools that sat on top but by the data architecture that sat underneath everything. Not the campaigns. Not the channels. Not the org chart. The architecture—because the architecture determines what's possible, and what's possible determines how you go to market.

The history of how companies collect, organize, and activate data is not just a technology story — it's the story of what becomes possible when the architecture changes. Every time the data foundation shifts, the go-to-market motion shifts with it.

"We have more data than any company in our industry. And we have less insight than a startup with a Snowflake account and two data engineers."

Figure 1: The Three Eras of Enterprise Data

Three eras, defined not by the marketing tools on top but by the data architecture underneath. The architecture determines what's possible, and what's possible determines how you go to market.

Dimension	Era 1: Warehouse 1990s–2010s	Era 2: Data Lake 2010s–2020s	Era 3: Data Fabric 2020s–present
The tagline	"A DMP knew your audience."	"A CDP knows your customer."	"A data fabric knows your business."
Architecture	Data warehouse (SAP BW, Teradata, Oracle). ETL: move data to analytics. Batch overnight.	Data lakes, lakehouses, CDPs. Fragmented — CRM, ERP, supply chain in separate systems. Unification problem.	Business data fabric (SAP BDC). Bring intelligence to the data. Semantic layer preserves context. Zero-copy exchange.
Data paradigm	Structured. Predefined questions. Periodic reports — monthly, quarterly, annually.	Unified customer view. Real-time signals. Identity resolution across systems.	Federated. Contextual. Agents reason across operational data with semantic understanding intact.
Who uses it	Small priesthood of analysts and finance. Everyone else gets dashboards or slide decks.	Broader access. Marketers, product teams, CX. Self-serve analytics begins.	AI agents as primary consumers. Humans direct; agents execute and synthesize.
PMM job	Static deliverables: positioning docs, messaging frameworks, campaign briefs. Quarterly cadence.	Expanded: lifecycle marketing, journey mapping, win/loss, persona development, data-driven pricing.	Redefined: from executor who produces artifacts to strategist who directs agents producing artifacts.
PMM skill set	Communications-heavy. Storytelling + enough technical fluency to be credible.	Data fluency. Credible with a CTO on data integration, explain it to a CMO in business terms.	Systems orchestration. Which parts automate, which need judgment, how to design the handoff.

The shift: The earlier transitions asked PMMs to expand their skill set. The fabric-to-agent shift asks PMMs to redefine their role entirely — from executors who produce artifacts to strategists who direct agents that produce at a pace no individual human could match.

Figure 1. The Three Eras of Enterprise Data. Each era reshaped not just technology stacks but organizational talent models. The companies that recognize the shift early build advantages that compound for years.

We think of this history in three eras, defined not by the marketing tools that sat on top but by the data architecture that sat underneath everything.

Figure 1: The Three Eras of Enterprise Data

Era One: The Data Warehouse and the Age of Reports

For most of the 1990s and 2000s, the center of gravity for enterprise data was the data warehouse. SAP BW, Teradata, Oracle—these were the

systems of record for analytics, and they operated on a simple premise: extract data from your operational systems, transform it into a structured model, and load it into a warehouse where analysts could run reports. ETL—extract, transform, load—was not just a technical process; it was a philosophy. You moved the data to the analytics. You waited for the batch job to run overnight. And in the morning, you had your numbers.

The data warehouse era was, in hindsight, wonderfully contained. The data was structured. The questions were predefined. The reports were periodic—monthly, quarterly, annually. And the people who consumed the data were a small priesthood of analysts and finance teams who knew how to write queries and build Excel models. The rest of the organization got dashboards, if they were lucky, or slide decks summarizing what the dashboards said.

For product marketers, this era defined the job in ways that persist to this day. The PMM's deliverables were fundamentally static: positioning documents, messaging frameworks, campaign briefs, competitive battlecards that were accurate on the day they were written and started decaying the moment they were saved. The Pragmatic Institute built its framework during this era, and the framework reflects the cadence of warehouse-era data—quarterly reviews, annual planning cycles, periodic competitive updates. The data moved slowly, so the marketing moved slowly, and that was fine because the buyers moved slowly too.

Chris O'Hara lived the marketing side of this era at Krux, one of the pioneering data management platforms that Salesforce later acquired. Even in the adjacent world of advertising data, the same warehouse-era assumptions held: you collected audience data in batches, segmented it into taxonomies, and activated it against campaigns that ran for weeks or months. The DMP—data management platform—was the warehouse paradigm applied to marketing, and it worked until it didn't.

Era Two: The Data Lake and the Age of Unification

The second era arrived not as a single product launch but as a slow-motion crisis of fragmentation. Enterprises woke up one morning—metaphorically, though some of them describe it exactly that way—and realized their data was everywhere. The warehouse still existed, but now there was also a data lake on AWS. A Snowflake instance that the analytics team had spun up. A Databricks workspace for the data science team. Half a dozen SaaS applications, each with its own data model and its own API. CRM over here, ERP over there, supply chain data in a third system, marketing data in a fourth. The dream of a single source of truth had shattered into a hundred sources of partial truth, none of them quite agreeing with the others.

This was the era that produced the customer data platform—the CDP—which Chris and Martin Kihn wrote the book on (literally: *Customer Data Platforms*, Wiley, 2021). The CDP was an attempt to solve the unification

problem for marketing: take all the customer data from all the systems, resolve the identities, and create a single customer view that marketers could actually use. The tagline that emerged from that work captured the generational shift: "A DMP knew your audience. A CDP knows your customer."

But the unification problem was never just a marketing problem. It was an enterprise problem. Every function—finance, supply chain, HR, procurement—was drowning in the same fragmentation. And the solutions that emerged reflected this: data lakehouses that tried to combine the structure of warehouses with the flexibility of lakes, data mesh architectures that tried to decentralize ownership while maintaining governance, and—most importantly for our story—the concept of a business data fabric that could connect data across systems without requiring you to move it all to one place.

This is where SAP's journey becomes central to the story, because SAP sits at the intersection of the problem in a way no other company does. SAP runs the operational backbone of the world's largest enterprises—the ERP, the supply chain, the financials, the HR systems. When we launched SAP Datasphere in 2023, the product strategy was built around a recognition that had been forming for years: customers didn't need another place to put their data. They needed a way to make their data meaningful wherever it already lived. Datasphere introduced the business data fabric—a semantic layer that preserved the business context of SAP data and connected it to non-SAP sources without forcing costly extraction and replication. It was, in retrospect, the bridge between the unification era and what came next.

For product marketers, the unification era expanded the job dramatically. The PMM was no longer just writing positioning and campaign briefs. Now the job included lifecycle marketing, personalization strategy, customer journey mapping, win/loss analysis, and a whole set of activities that required understanding individual customers and their relationship with the product. The Pragmatic Institute's framework grew to accommodate this—more activities around customer insights, persona development, data-driven pricing. The PMM became less of a campaign strategist and more of a customer strategist.

Era Three: The Data Fabric and the Age of Agents

Which brings us to now. And the line that captures this era: a DMP knew your audience. A CDP knew your customer. A data fabric—powered by agentic AI—knows your business.

In February 2025, we launched SAP Business Data Cloud—a fully managed SaaS platform that unified SAP Datasphere, SAP Analytics Cloud, and SAP Business Warehouse with an embedded Databricks partnership, all built around a shared object store and zero-copy data exchange. The product strategy behind BDC was something we had been building toward

for years: stop asking customers to move their data to the analytics, and instead bring the intelligence to the data. Build semantic data products that carry their business context with them—so that "Customer" in SuccessFactors and "Customer" in S/4HANA don't get flattened into the same field by a junior data engineer at 2 AM. Federate the metadata. Keep the data where it lives. And make the whole thing the foundation for AI that actually understands the business it's operating in.

We need to be careful here, because "agentic" is rapidly becoming one of those words that means everything and nothing, like "digital transformation" or "cloud-native." When we say agentic, we mean something specific: software systems that can perceive their environment, make decisions, and take actions without waiting for a human to tell them what to do at each step. Not chatbots. Not copilots that suggest the next sentence in your email. Agents that monitor your competitive landscape while you sleep. Agents that adjust your pricing recommendations based on real-time market signals. Agents that generate and test messaging variants based on what's resonating with different buyer segments. Agents that evaluate your demo environment and flag when it's showing a workflow that doesn't match the customer's use case.

The data fabric sits at the convergence of three things that didn't exist—or didn't exist at sufficient scale—five years ago. First, large language models that can reason about unstructured business context in ways that previous AI could not. Second, rich enterprise data fabrics—like SAP Business Data Cloud—that give those models access to the operational data (ERP, CRM, supply chain, finance) that actually runs the business, with the semantic context intact. Third, agent orchestration frameworks that allow multiple AI agents to collaborate on complex, multi-step workflows without constant human supervision. HANA Cloud's knowledge graph, its multi-modal compute engine spanning relational, document, graph, and knowledge graph processing—these are not features on a product roadmap. They are the infrastructure that makes enterprise-grade agentic AI possible.

The data fabric sits at the convergence of three things that didn't exist — or didn't exist at sufficient scale — five years ago. Together, they make enterprise-grade agentic AI possible.

Figure 2: The Agentic Convergence

Three capabilities converging to create the agentic era. None is sufficient alone. All three together change who does the work.

LARGE LANGUAGE MODELS	ENTERPRISE DATA FABRICS	AGENT ORCHESTRATION
<p>What's new</p> <p>Models that can reason about unstructured business context in ways previous AI could not. Not just pattern matching — contextual understanding.</p> <hr/> <p><i>An agent that reads a competitor earnings call transcript, extracts the three most strategically significant statements, and explains why they matter for your positioning.</i></p>	<p>What's new</p> <p>Rich operational data (ERP, CRM, supply chain, finance) with semantic context intact. "Customer" in SuccessFactors and "Customer" in S/4HANA don't get flattened.</p> <hr/> <p><i>SAP Business Data Cloud: federated metadata, zero-copy exchange, semantic data products that carry their business context with them.</i></p>	<p>What's new</p> <p>Frameworks that allow multiple AI agents to collaborate on complex, multi-step workflows without constant human supervision.</p> <hr/> <p><i>Agents that monitor your competitive landscape while you sleep, adjust pricing recommendations on market signals, generate and test messaging variants.</i></p>



THE AGENTIC ERA

Software systems that can perceive their environment, make decisions, and take actions without waiting for a human at each step. Not chatbots. Not copilots. Agents.

A DMP knew your audience. A CDP knew your customer. A data fabric knows your business.

What this changes for PMMs: The earlier shifts asked you to expand your skill set. This shift asks you to redefine your role — from executor who produces artifacts to strategist who directs agents that produce artifacts at a pace and quality no individual human could match.

Figure 2. The Agentic Convergence. LLMs provide reasoning. Data fabrics provide context. Agent orchestration provides autonomy. Together, they change not what data you have access to, but who does the work.

"Agentic" is rapidly becoming one of those words that means everything and nothing. When we say agentic, we mean something specific: software systems that can perceive, decide, and act without waiting for a human at each step.

Figure 2: The Agentic Convergence

For product marketers, this changes the game in a way that's different from the warehouse-to-lake transition. That transition changed what data you had access to and what you could do with it. This transition changes who does the work. The earlier shifts asked PMMs to expand their skill set. The fabric-to-agent shift asks PMMs to redefine their role entirely—from executors who produce artifacts to strategists who direct agents that produce artifacts at a pace and quality that no individual human could match.

That is an uncomfortable sentence to write, and probably an uncomfortable one to read. But we've watched enough hype cycles to know that the uncomfortable truths are the ones worth paying attention to early.

Each data era reshaped not just the technology stack but the PMM's role, deliverables, and required skill set. The progression is clear — and the third shift is the most disruptive.

Figure 3: The PMM Role Shift Across Eras

How the PMM job changed with each architectural shift. The first two transitions expanded the job. The third redefines it.

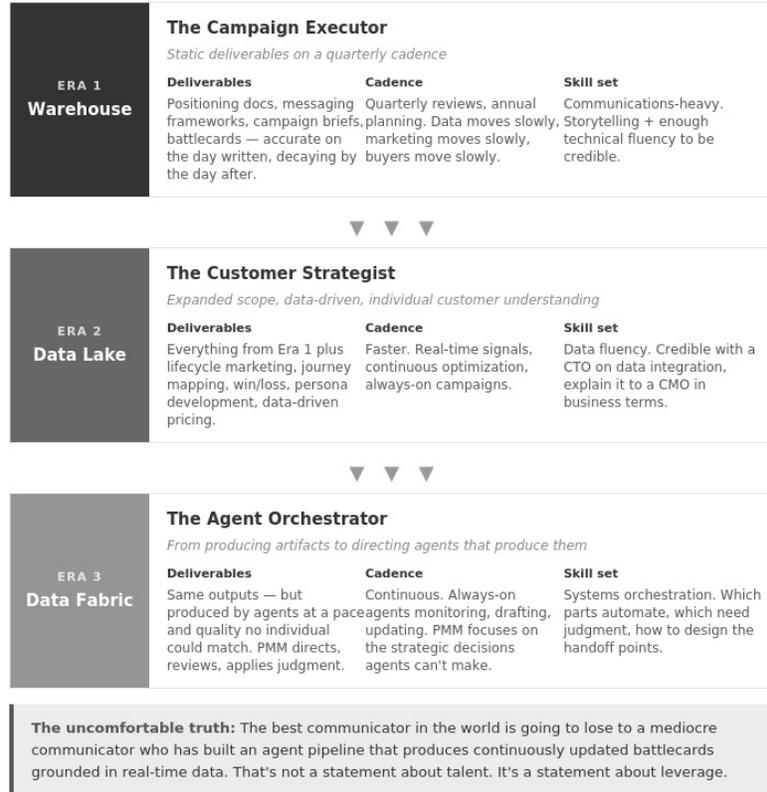


Figure 3. The PMM Role Shift Across Eras. Era 1 expanded to Era 2. Era 2 expanded to Era 3. But Era 3 is different — it doesn't just add activities. It changes who does them.

The fabric-to-agent shift asks PMMs to redefine their role entirely — from executors who produce artifacts to strategists who direct agents that produce artifacts at a pace and quality that no individual human could match.

Figure 3: The PMM Role Shift Across Eras

• • •

The rest of this book is organized around a simple question: for each major activity in the product marketing job, what changes when agents become part of the team? Not what changes in theory. What's already changing, right now, for the PMMs who are paying attention.

But before we get there, we need to do something that we think nobody has done honestly yet: take the canonical framework for product marketing and hold it up to the light. Chapter 2 is about the Pragmatic Institute's thirty-seven activities, and what happens when you score each one for AI disruption potential. The results are sobering—and, if you're inclined

toward optimism, genuinely exciting.

The CMO's perspective maps each era to an organizational architecture and hiring profile. The progression has moved from communications to data fluency to systems orchestration — and leverage, not effort, is the new differentiator.

Figure 4: The Leverage Equation

What CMOs need from PMMs in each era — and why the agentic era demands a fundamentally different talent model.

Dimension	Era 1: Warehouse	Era 2: Data Lake	Era 3: Data Fabric
Hiring profile	Communications-heavy. Storytelling plus enough technical fluency to be credible.	Data fluency. Credible with a CTO on integration, explain it to a CMO in business terms.	Systems thinker. Can look at a complex workflow and figure out what to automate vs. what needs judgment.
Differentiator	Quality of writing and messaging. The best communicator wins.	Depth of customer understanding. The most data-fluent PMM wins.	Leverage. The PMM who builds agent pipelines beats the one who manually produces artifacts.
Org design	Well-defined roles: messaging, positioning, campaign support. Functional silos.	Broader roles. PMMs needed across lifecycle, personas, pricing, analytics. Cross-functional.	Orchestration roles. PMMs design workflows, direct agents, apply judgment at handoff points.
Capacity model	Linear. Output = headcount x hours. More people = more output.	Partially leveraged. Better data = better targeting = more efficient campaigns.	Exponential. One PMM with agent pipelines outproduces a team without them.
Risk of being late	Slow decay. Competitors with better messaging gradually win positioning.	Faster erosion. Data-driven competitors personalize while you're still in quarterly cycles.	Structural gap. Competitors with agent leverage produce 10x while you produce 1x. The gap compounds.

***Leverage, not effort, is the new differentiator.
Building agent pipelines beats manually producing artifacts.***

For the book ahead: Every remaining chapter asks a version of the same question — for each major activity in the PMM job, what changes when agents become part of the team? Not in theory. What's already changing, right now, for the PMMs who are paying attention.

Figure 4. The Leverage Equation. Each era shifted the talent model. In the agentic era, the differentiator isn't the quality of what you produce — it's the leverage of how you produce it. The gap between agent-enabled PMMs and manual-execution PMMs compounds over time.

But before we get there, we need to do something nobody has done honestly yet: take the canonical framework for product marketing and hold it up to the light.

Figure 4: The Leverage Equation

THE CMO PERSPECTIVE

The three-era framework maps not just to data architecture but to organizational architecture. In the warehouse era, the PMM job was well-defined: messaging, positioning, campaign support. The skill set was

communications-heavy—storytelling plus enough technical fluency to be credible. In the unification era, the profile shifted. You needed people who could have a credible conversation with a CTO about data integration and identity resolution and then explain it to a CMO in business terms. The job got harder, the skill set got broader, but it was still fundamentally a human-execution job.

In the agentic era, the hiring profile changes again. What we need now are people who can orchestrate—who can look at a complex workflow like building a competitive battlecard and figure out which parts should be automated, which parts need human judgment, and how to design the handoff points. That's a systems thinking skill, not a communications skill. The best communicator in the world is going to lose to a mediocre communicator who has built an agent pipeline that produces continuously updated battlecards grounded in real-time data. That's not a statement about talent. It's a statement about leverage.

KEY TAKEAWAYS

- Each data era—warehouse, lake, fabric—reshaped not just technology stacks but organizational talent models.
- The PMM skill progression has moved from communications to data fluency to systems orchestration.
- Leverage, not effort, is the new differentiator: building agent pipelines beats manually producing artifacts.
- The agentic era demands that every PMM understand where human judgment ends and automation begins.

The Pragmatic Framework Meets the Machine

If you've been through Pragmatic Institute training—and there's a decent chance you have, since they've certified over 250,000 professionals—you know the Framework. Thirty-seven boxes arranged in seven categories: Market, Focus, Business, Planning, Programs, Enablement, and Support. Each box represents an activity that a product team needs to execute well to bring a product to market successfully. It's been the closest thing the product management and product marketing disciplines have to a shared operating system for three decades.

We've stared at that framework a lot over the past year. Not because we were studying for a certification—we were trying to answer a question that we couldn't find anyone else asking in a rigorous way: what happens to each of those thirty-seven boxes when AI agents become part of the team? Not in the abstract, feel-good, "AI will augment human creativity" sense. Concretely. Activity by activity. Which ones does an agent do better than a human today? Which ones does an agent make dramatically faster without sacrificing quality? And which ones remain stubbornly, irreducibly human—the ones where judgment, taste, political awareness, and lived experience are the actual product?

The exercise turned out to be both clarifying and unsettling in roughly equal measure. We'll walk you through the methodology, the results, and what we think it means for every PMM who's trying to figure out where to invest their time and energy.

The Scoring Methodology

We scored each of the thirty-seven Pragmatic activities across three dimensions. If you're the kind of person who likes to see the math—and if you're a PMM, you probably are—we built an interactive version of this analysis on futureofpmm.com where you can adjust the weightings and see how the scores shift. But for the purposes of this chapter, here's the framework.

the threshold for perfection. It is the threshold for 'good enough to act on,' which is the threshold that matters for procurement decisions, competitive intelligence, and content production.

Figure 1: AI Replacement Probability by Activity

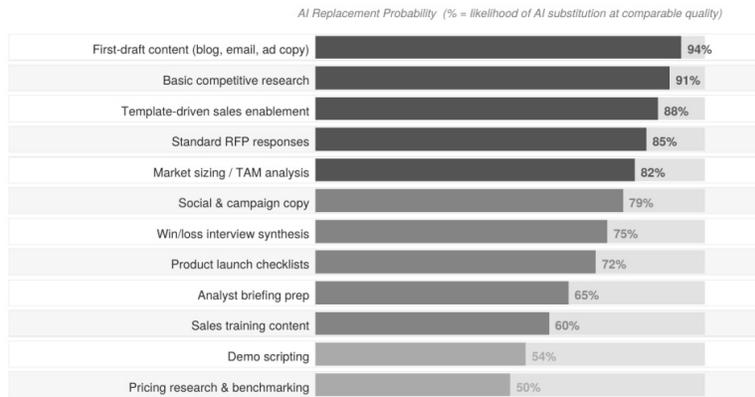


Figure 1: AI Replacement Probability by Activity

The first dimension is **automation potential**: on a scale of 1 to 5, how much of this activity can an AI agent perform today, in early 2026, with commercially available tools? A score of 1 means the activity is almost entirely human. A score of 5 means an agent can do it end-to-end with minimal human oversight. We were deliberately conservative here—we scored based on what we've seen work in practice, not what a vendor's demo deck promises.

The second dimension is **augmentation value**: on a scale of 1 to 5, how much better does a human-plus-agent team perform than a human alone? This is different from automation. Some activities can't be fully automated but become dramatically more effective when an agent handles the research, synthesis, or first-draft generation. A score of 1 means the agent adds marginal value. A score of 5 means the human-agent combination is transformatively better than the human solo.

The third dimension is **elevation opportunity**: on a scale of 1 to 5, does this activity become more strategically important when the routine work around it is automated? This is the dimension most people miss. When you free up the hours that used to go into gathering competitive data, the act of interpreting that data and making strategic decisions based on it doesn't just stay the same—it gets elevated. The PMM who used to spend 80% of their competitive intelligence time on collection and 20% on analysis can now flip that ratio. The analysis itself becomes more valuable because there's more of it, it's grounded in better data, and it's happening in something closer to real time.

• • •

Running this analysis produced three natural clusters that organize the rest of this book. We're going to call them what they are, without softening

the language, because PMMs deserve honesty about what's coming.

The Activity Matrix

Replacement probability is one axis. The other is strategic value — how much does this activity actually move the needle on the outcomes that matter to the business: revenue, win rate, market share, competitive position? Plotting the two axes produces a matrix that is more diagnostic than either axis alone.

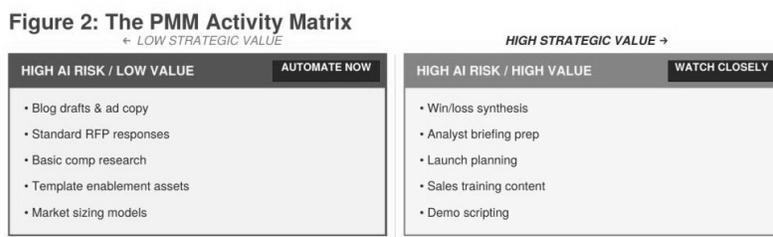


Figure 2: The PMM Activity Matrix

Cluster One: Automated

These are activities where automation potential scored 4 or 5—meaning that AI agents can perform the majority of the work today, with human oversight reduced to quality checks and strategic direction. This cluster includes activities that are fundamentally about data gathering, pattern matching, and content generation from structured inputs.

Take **Market Sizing**. The traditional approach: a PMM spends two to three weeks pulling data from Gartner, IDC, Statista, company filings, and industry reports. They build a bottom-up model in Excel. They triangulate against a top-down estimate. They produce a slide deck that says the addressable market is \$4.2 billion growing at 12% CAGR. It's essential work—you can't build a business case without it—but it's also the kind of work that an agent with access to the right data sources can do in an afternoon. Not approximately. Actually well. The agent reads the same reports, builds the same models, triangulates the same way. The human's job shifts from doing the analysis to validating it and challenging the assumptions.

Or take **Content Creation**—specifically, the production of derivative content. Once you have a core messaging framework and a strong point of view, the mechanical act of turning that into blog posts, social copy, email sequences, slide decks, and sales one-pagers is straightforward generation work. Chris tested this extensively for futureofpmm.com, running the same content brief through Claude, ChatGPT, and Prezent. The results weren't perfect—they never are—but the first-draft quality was good enough that the human's job became editing and elevation rather than creation from scratch. That's a 70% time savings on the production side, conservatively.

The Two Halves of PMM

The matrix simplifies cleanly into two halves. There is the commodity work — the content that gets produced, the research that gets synthesized, the templates that get filled — where AI is already competitive with human execution and will only become more so. And there is the irreplaceable work — the narrative that gets architected, the category that gets created, the executive conversation that shifts a strategic decision — where human judgment is still the core of the value.

Figure 3: The Two Halves of PMM

COMMODITY WORK — Going to Near-Zero Cost	IRREPLACEABLE WORK — Growing in Value
First-draft content <small>Blog posts, emails, ad copy, social — AI produces 80%+ quality d</small>	Narrative architecture <small>Making a technical product feel like an obvious choice. AI can draft; it c</small>
Basic research	Competitive strategy

Figure 3: The Two Halves of PMM

Other activities that cluster here: **Event Support** operations (logistics, asset tracking, coordination), much of the **Sales Tools** production pipeline, routine aspects of the **Marketing Plan** (channel mapping, budget allocation against established models), and the mechanical dimensions of **Competitive Analysis**—the gathering part, not the interpretation part. We'll come back to that distinction in Chapter 3, because it matters enormously.

Here's the uncomfortable truth about Cluster One: these are the activities that many PMMs spend the majority of their time on. If you're a PMM whose primary value is producing artifacts—the battlecard, the one-pager, the launch email sequence, the competitive slide—the automation wave is not your friend. Those artifacts are going to be produced by agents at a fraction of the cost and a multiple of the speed. The question isn't whether this happens. It's whether you're the person directing the agents or the person being replaced by them.

Cluster Two: Augmented

These activities scored high on augmentation value (4 or 5) but moderate on automation potential (2 or 3). They're activities where the human remains essential but where the human-agent team is so much better than the human alone that operating without agents becomes a competitive disadvantage.

Positioning is the canonical example. No agent is going to write your positioning from scratch—not because the technology can't generate positioning statements (it can, all day long) but because positioning is fundamentally a strategic choice about what you want to be known for, who you're willing to lose, and how you want to show up relative to alternatives. Those are judgment calls that require understanding the market, the product, the competitive landscape, the sales team's capabilities, and the company's strategic ambitions all at once. An agent

doesn't have that context. But an agent can synthesize the inputs that inform those judgment calls—win/loss data, competitive positioning shifts, analyst commentary, customer verbatims—at a speed and comprehensiveness that no human can match. The PMM who makes positioning decisions based on an agent-synthesized intelligence brief is going to make better decisions than the PMM who makes them based on gut instinct and a quarterly competitive review.

Win/Loss Analysis falls squarely in this cluster. The traditional process—interview the sales rep, maybe interview the customer if they'll talk to you, write up the findings, identify patterns, present to leadership—is time-intensive and limited by sample size. A PMM running win/loss analysis might cover twenty deals in a quarter. An agent-augmented PMM can process hundreds of deal records, pull themes from call transcripts, cross-reference against CRM data, and surface patterns that would take a human months to identify. The human still needs to interpret those patterns, decide which ones are actionable, and figure out the organizational response. But the input quality is transformatively better.

Buyer Personas sit here too. The research that feeds persona development—customer interviews, survey data, behavioral analytics, market segmentation—is exactly the kind of synthesis work that agents excel at. But the act of making a persona useful—deciding which attributes matter for positioning, which pain points are real versus performative, which objections are deal-breakers versus negotiating tactics—requires the kind of contextual judgment that comes from having sat across the table from actual buyers. An agent can build you a richly detailed persona document in an hour. Whether that document reflects reality depends on the human who directed the research and validated the output.

Other Cluster Two activities: the **Buying Process** mapping, **Use Case** development, **Pricing** strategy, **Product Roadmap** input, and the strategic layer of the **Go-to-Market Plan**. In each case, the agent accelerates the inputs and the human elevates the decision.

Cluster Three: Elevated

This is where it gets interesting—and where the career implications are most significant. Cluster Three activities scored high on elevation opportunity (4 or 5) and low on automation potential (1 or 2). These are the activities that become more important, more central, and more career-defining in the agentic era precisely because they are the ones that machines can't do.

Stakeholder Management. Every PMM knows that the hardest part of the job isn't writing the positioning document—it's getting twelve people to agree on it. It's navigating the VP of Sales who wants feature-led messaging, the CPO who wants vision-led messaging, the CFO who wants ROI-led messaging, and the CEO who wants all three on a single slide. No agent is going to sit in that room and read the body language and know

that the reason the VP of Sales is pushing back isn't because they disagree with the positioning but because they're worried about quota attainment and they need the messaging to validate the deal structure they've been pitching for the last quarter. That's organizational intelligence. It's learned, not computed.

Thought Leadership—real thought leadership, not the commodity-content kind—lands here. Chris wrote about this at length on [futureofpmm.com](#): the paradox is that AI makes it trivially easy to produce content that looks like thought leadership but contains no actual thinking. The market is about to be flooded with competent, bland, algorithmically generated "insights." The PMM who can cut through that noise with a genuine perspective—earned through experience, grounded in specific stories, delivered with a voice that is recognizably human—is going to be more valuable than ever. Not because thought leadership is new, but because the signal-to-noise ratio is about to collapse, and being the signal is a differentiator.

The most uncomfortable finding in the analysis is not the replacement probability scores. It is the correlation between career stage and exposure. The PMMs who are most at risk of AI substitution are, by and large, the PMMs who have spent the most time building expertise in exactly the activities that AI is replacing. And because that expertise feels hard-won and valuable — because it took years to develop and clients have paid for it — those PMMs are the least likely to perceive the threat.

Figure 4: The Vulnerability Paradox

The Strategic PMM	The Mid-Career PMM	The Execution PMM
Typical tenure 5–15 yrs	Typical tenure 2–5 yrs	Typical tenure 0–3 yrs
Time on commodity work ~20% on commodity tasks	Time on commodity work ~50% on commodity tasks	Time on commodity work ~75% on commodity tasks
Has built expertise in narrative, positioning, and e	Splitting time between strategic and execution wo	Built career on volume and execution speed. Most of
Uses AI to accelerate research and drafts. Netbo	is open — but it requires deliberate investment in	their daily work sits in the high-replacement zone. And

Figure 4: The Vulnerability Paradox

Customer Empathy is the one that surprises people, but it shouldn't. An agent can analyze ten thousand support tickets and tell you that 23% mention onboarding friction. A human can sit across from a customer, watch their face when they describe trying to get their team to adopt a new tool, and understand that the real problem isn't the onboarding UX—it's that the customer doesn't have executive sponsorship and they're trying to drive adoption bottom-up in a top-down culture. That's a completely different product marketing challenge than "fix the onboarding flow," and no amount of ticket analysis will surface it.

Other Cluster Three activities: **Distinctive Competence** identification—figuring out what your company is uniquely good at and why it matters; the strategic dimension of **Channel Management**—choosing the right partners and managing those relationships; and **Innovation Games**—the facilitated exercises where you discover unmet needs that customers can't articulate because they don't know they have them yet. In every case, the

activity requires something an agent doesn't have: a theory of mind about other humans, a sense of organizational dynamics, and the ability to make judgment calls under ambiguity.

• • •

The 10x Map

Here's what the three clusters mean in practical terms for how you spend your time.

If you're a PMM who currently spends 60% of your week on Cluster One activities—producing artifacts, gathering data, updating decks, coordinating logistics—you have a choice. You can keep doing those things the way you've always done them, in which case you're going to find yourself increasingly redundant as agents get better and as your more forward-leaning colleagues demonstrate what's possible. Or you can invest the next six months in building agent-powered workflows for those activities, reclaim those hours, and reinvest them in Cluster Two and Cluster Three work.

The math isn't complicated. A PMM who automates 20 hours a week of Cluster One work and redirects that time into deeper competitive analysis, more thoughtful positioning, better stakeholder management, and genuine thought leadership is going to outperform a PMM who's still manually updating battlecards and building market sizing models in Excel. Not by 20%. By a factor that justifies the title of this book.

That's not a threat. It's an opportunity—arguably the biggest opportunity in the history of the PMM discipline. Product marketing has always suffered from being spread too thin. The Pragmatic Framework has thirty-seven activities because the job really does encompass all of that. No single person can do all thirty-seven well. The agentic era doesn't shrink the list—it makes it possible, for the first time, for an individual PMM to meaningfully cover the full scope while going deep on the activities that matter most.

That is what "10x yourself" means. Not working ten times harder or ten times faster. Operating at a fundamentally different level of strategic impact because you've offloaded the mechanical work to machines that are built for it, and you're finally free to do the work that only a human can do.

• • •

The rest of Part II takes this framework and applies it to each major cluster of PMM activity. Chapter by chapter, we'll walk through what's changing, what the best PMMs are doing about it, and what you can start doing Monday morning. Each chapter maps back to specific Pragmatic activities,

so if you want to trace the thread from the Framework to the agentic reality, you can.

But first, it's worth looking at what this analysis means from the other side of the hiring table—the executive perspective on what actually separates a good PMM from a great one.

THE CMO PERSPECTIVE

The three-cluster analysis isn't just a useful framework—it's how CMOs are already making team decisions, whether they had a name for it or not. Over the past year, evaluating PMM performance has increasingly meant weighting Cluster Three contributions—stakeholder navigation, strategic judgment, genuine customer insight—far more heavily than Cluster One output. A great battlecard used to be impressive. Now the question is: did you build a system that produces great battlecards continuously?

That sounds harsh, and it deserves nuance. The artifact still matters. But the artifact is the output, not the value. The value is the insight that went into it, the strategic choice it represents, and the business outcome it drives. In the agentic era, a well-prompted agent can produce the artifact. What it can't produce is the judgment that decides what artifact to build, who it's for, and why it matters right now. When hiring, the most revealing question isn't "show me what you've produced"—it's "tell me how you work."

KEY TAKEAWAYS

- The Pragmatic Framework's 37 activities cluster into three tiers: automatable, augmentable, and irreducibly human.
- CMOs are already weighting Cluster Three skills—judgment, navigation, insight—over artifact production.
- The question has shifted from "can you produce a great deliverable?" to "can you build a system that produces great deliverables continuously?"
- Interview signals that matter: how candidates describe their process and their relationship with AI tools.

Market Intelligence and Competitive Strategy

Pragmatic Remix: Competitive Landscape • Win/Loss Analysis • Market Sizing • Market Problems

Soo Shim had a problem that every competitive intelligence PMM recognizes: she was drowning in data and starving for insight. As Senior Director of Global Competitive Intelligence at SAP, Soo was responsible for tracking dozens of competitors across multiple product lines—enterprise data clouds, application platforms, analytics suites, planning tools. Her team produced battlecards, ran win/loss programs, briefed sales teams, and fed intelligence into product roadmap discussions. By any traditional measure, they were good at it.

But "good at it" had a ceiling, and Soo could feel it. The battlecards were accurate on the day they were published and slightly stale a week later. The win/loss analyses surfaced useful patterns but took weeks to produce and covered a fraction of the deals they should have. The competitive landscape was shifting faster than any team of humans could track—a competitor would announce an acquisition on a Tuesday, reposition their narrative on a Wednesday, and by Thursday their sales reps were in the field with new talking points that Soo's team hadn't seen yet.

So Soo built BattleCoach.

The freshness timestamp is not a vanity metric. It is the single most important attribute the Living Battlecard gains over the quarterly battlecard — because it gives the rep permission to trust the document. A battlecard with no date is a document with an unknown expiration.

Figure 4: The BattleCoach Flywheel

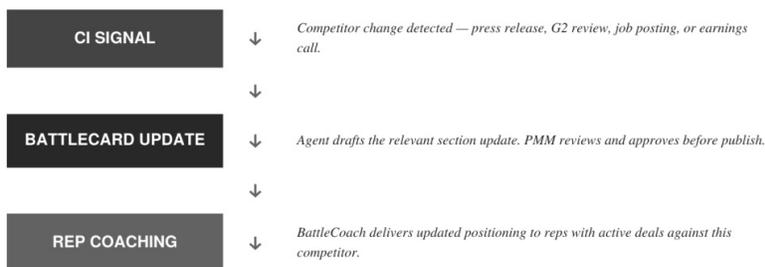


Figure 1: The BattleCoach Flywheel

BattleCoach isn't a monitoring dashboard or an alerting tool—though it

does both of those things. It's an AI-powered competitive coaching platform that transforms static competitive content into dynamic, personalized training for sales reps. The system ingests competitive data continuously—analyst reports, earnings calls, product announcements, pricing changes, customer reviews, social media signals—and synthesizes it into deal-specific intelligence. A sales rep preparing for a call against Snowflake doesn't open a PDF battlecard; they ask BattleCoach what's changed in Snowflake's positioning this week, what objections they're likely to face based on the prospect's industry and company size, and how to reframe the conversation around SAP's differentiated capabilities.

The gap between what Soo's team was doing before BattleCoach and what they're doing with it isn't incremental. It's categorical. The team didn't get faster at producing battlecards. They stopped producing static battlecards entirely. The artifact was replaced by a living system.

• • •

Soo's story illustrates the transformation happening across competitive intelligence—and, more broadly, across everything the Pragmatic Framework groups under "Market." The activities in this cluster are the ones where the automation-plus-augmentation potential is highest, which means they're also the ones where the gap between agent-powered PMMs and traditional PMMs is widening fastest.

Let us walk through what's changing, activity by activity.

Competitive Landscape: From Periodic Review to Persistent Monitoring

The traditional competitive intelligence cycle looked something like this: once a quarter, a PMM or a dedicated CI analyst would update the competitive landscape assessment. They'd review recent announcements, check analyst reports, talk to sales about what they were hearing in deals, and produce a document—usually a slide deck or a matrix—that summarized where each competitor stood on key dimensions. That document would get presented to leadership, distributed to sales, and then gradually decay in accuracy until the next quarterly refresh.

We've lived this cycle for years, and we can tell you exactly what was wrong with it: the cadence didn't match the market. Competitive landscapes don't shift on a quarterly schedule. They shift when Databricks announces a partnership with Salesforce. When Snowflake drops pricing on a key workload. When a competitor's CTO publishes a blog post that subtly repositions their platform away from data warehousing and toward AI orchestration. Those signals don't wait for your quarterly review.

reps into deals with positioning that the market has already moved past.

The solution is not to update battlecards more frequently. That is a headcount problem disguised as a process problem. The solution is to change the architecture: from a write-once, publish-quarterly system to a living system that continuously monitors signals, drafts updates, and delivers intel in context. AI makes this architecture practical for the first time. The question is whether the PMM organization is willing to build it.

Figure 1: The CI Maturity Model

STAGE 1 REACTIVE	STAGE 2 PERIODIC	STAGE 3 MONITORED	STAGE 4 LIVING SYSTEM
<i>No system</i>	<i>Quarterly cycle</i>	<i>Alerts + triggers</i>	<i>Always-on agents</i>
CI happens when a rep loses a deal and asks what happened. Battlecards don't exist or are 2+ years old.	A dedicated PMM updates battlecards once a quarter. Good coverage at point of publication. Stale within weeks.	Automated alerts fire on key competitor signals. PMM reviews and updates within days. Battlecards carry	Agents continuously monitor signals, draft updates, and push changes to reps in context — in Slack, the

Figure 2: The CI Maturity Model

Chris built an agent-powered CI monitoring system—he wrote about it in detail on futureofpmm.com—that tracks positioning shifts, pricing changes, product announcements, executive messaging, analyst coverage, and social sentiment across more than fifteen competitors. The system runs continuously. It doesn't produce a quarterly deck; it produces a daily intelligence brief that highlights what's changed, what it means, and what—if anything—requires a response.

The architectural pattern is straightforward, and we want to lay it out because we think every PMM should understand how this works, even if they're not the one building it. You start with **data ingestion**: RSS feeds from competitor blogs and newsrooms, monitoring of earnings call transcripts, tracking of analyst report publications, social listening on key executive accounts, and alerts on product documentation changes. Those feeds go into a **processing layer** where an LLM synthesizes the raw signals into structured intelligence—categorizing each item by competitor, topic (pricing, positioning, product, people, partnerships), and significance level. The structured intelligence flows into a **synthesis layer** where the agent compares today's signals against the established competitive baseline and identifies meaningful shifts. And the **output layer** delivers the intelligence to the right people in the right format—a Slack alert for urgent signals, a daily brief for the CI team, a weekly summary for leadership, and deal-specific intelligence that plugs into the sales enablement workflow.

in the signals they already use, in the context of the deals they are currently tracking.

The move from Stage 2 to Stage 4 is not primarily a resource question. A well-configured agent stack can monitor more signals more consistently than any team of analysts. The investment is in architecture: building the signal layer, configuring the analysis pipeline, and integrating the distribution layer with the CRM and sales tooling.

Figure 2: The Living CI Architecture

SIGNAL LAYER — What gets monitored		
Press releases & company blogs	G2 / TrustRadius reviews	Job postings & hiring patterns
Earnings calls & SEC filings	Social channels & community forums	Win/loss call transcripts
ANALYSIS LAYER — What agents do with signals		
Pattern detection across all sources	Change-delta vs. last snapshot	Sentiment shift in review scores
Draft battlecard update sections	Flag high-priority alerts for PMM	Link signal to open deals in CRM
DISTRIBUTION LAYER — How intel reaches reps		

Figure 3: The Living CI Architecture

Here's what we want you to notice about that architecture: none of the components are technically exotic. RSS feeds have been around since the early 2000s. Monitoring tools existed before AI. The difference is the synthesis layer. Before LLMs, you could gather all of those signals, but making sense of them—connecting a pricing change to a positioning shift to an analyst comment to a product announcement and understanding what the combined pattern means—required a human analyst with deep domain knowledge and a lot of time. Now an agent can produce the first-pass synthesis in minutes, and the human analyst's job shifts from "what happened?" to "what should we do about it?"

Win/Loss Analysis: From Sample to Census

Win/loss has always been one of the highest-value, lowest-frequency activities in the PMM toolkit. The Pragmatic Framework includes it because it's essential: understanding why you won and why you lost is the most direct input into positioning, messaging, competitive strategy, and product roadmap priorities. But the traditional approach—structured interviews with buyers and sales reps, typically covering ten to twenty deals per quarter—suffers from sample size problems. Twenty deals out of a pipeline of two hundred is a 10% coverage rate, and the selection is usually biased toward deals the sales team is willing to discuss, which skews the data.

Agent-augmented win/loss changes the coverage equation dramatically. An agent can process every deal in your CRM—not just the ones you have time to interview. It reads the call notes, the email threads, the Gong or Chorus transcripts, the opportunity field data. It identifies patterns that a human analyst reviewing twenty deals would never see: that deals

involving a specific competitor have a 40% longer sales cycle when the technical evaluator is from a particular department, or that your win rate drops by fifteen points when the prospect has previously used a specific competitive product, or that pricing objections correlate with deal size in a non-linear way that your current discounting model doesn't account for.

We want to be clear about what the agent does and doesn't do here. It doesn't replace the human win/loss interview. A conversation with a buyer who chose your competitor will always surface insights that no amount of CRM data mining can replicate—the hesitation in their voice when they talk about the other vendor's support experience, the offhand comment about how their CFO reacted to the pricing proposal, the admission that the deal was really decided by an internal champion who left the company before the contract was signed. Those are human-to-human insights. But the agent gives you something the interview can't: comprehensive pattern recognition across your entire deal population. The combination—agent for breadth, human for depth—is categorically better than either approach alone.

AEs running deals where data security is a key concern.

The Distribution Layer is where most CI systems are weakest. Intel that reaches the battlecard but does not reach the rep in context is only marginally better than intel that never leaves the PMM's Notion page. The Living CI Architecture distributes via push: Slack alerts to AEs with relevant active deals, BattleCoach prompts triggered by deal stage, weekly digests to CS and SE teams, and briefing documents for analyst conversations. PMMs own QA at the boundary between Analysis and Distribution — the human checkpoint before any drafted update reaches a rep.

Figure 3: The Battlecard Rebuild

DIMENSION	BEFORE: The Quarterly Battlecard	AFTER: The Living Battlecard
Format	Static Google Doc or PDF. No version control.	Structured data record in Notion or Confluence, updated via agent write-back.
Update cycle	Quarterly, when PMM finds the time. Often less.	Triggered by signal threshold — within 24–48 hours of meaningful change.
Content	Feature comparisons, talk tracks, objection	Same sections plus: freshness timestamp.

Figure 4: The Battlecard Rebuild

The Practitioner's Playbook: Getting Started with Agent-Powered CI

If you're a PMM reading this and thinking "this sounds great but I don't know where to start," here's what we'd recommend based on what we've seen work.

Start with the daily brief, not the monitoring system. The full architecture we described above is a multi-week build. But you can get 80% of the value in a day by setting up a simple workflow: every morning, prompt your LLM of choice with a structured request to synthesize the latest news, blog posts, and social mentions for your top three competitors. Feed it the raw URLs or use a tool like Perplexity that has built-in web access. Ask it to categorize what's changed into positioning, pricing,

product, and people dimensions. Ask it to flag anything that represents a meaningful shift from the competitor's established narrative. You'll be amazed at how much this surfaces that your current process misses.

Second, build the win/loss synthesis before you build the monitoring system. Export your last two quarters of closed-won and closed-lost deals from your CRM, including all the field data and any attached notes or transcripts. Run it through Claude or ChatGPT with a structured prompt that asks for pattern identification across win reasons, loss reasons, competitive presence, deal size, sales cycle length, and any other dimensions you track. The output won't be perfect—it can't be, because CRM data is messy and incomplete—but it will surface patterns you didn't know were there, and that's where the real value lives.

Third—and this is the one most people skip—build the distribution mechanism before you build the intelligence. The best competitive intelligence in the world is worthless if it lives in a Notion doc that nobody reads. Before you invest in gathering and synthesizing better intelligence, figure out how it's going to reach the people who need it: the sales rep five minutes before a call, the product manager deciding what to prioritize next quarter, the leadership team making pricing decisions. The delivery mechanism—Slack alerts, CRM integrations, deal-specific briefings—is what turns intelligence into action.

THE CMO PERSPECTIVE

Competitive intelligence in most organizations is underfunded, under-resourced, and underwhelming—often a one-person function or a responsibility shared across multiple PMMs who each give it ten percent of their attention. The result is predictable: CI is adequate but not excellent, and leadership doesn't invest more because they've never seen what excellent CI looks like. Agent-powered CI breaks that cycle because it changes the output so dramatically that leadership notices.

When our team demonstrated what BattleCoach could do—real-time, deal-specific competitive intelligence that sales reps actually used—the response wasn't "nice improvement." It was "how do we scale this across every product line?" That's the conversation every CI PMM wants to have with their leadership. Our advice: don't ask for permission to build this. Build it on a small scale—three competitors, one product line, a daily brief to your sales team—and let the results speak.

KEY TAKEAWAYS

- Agent-powered CI transforms a chronic underinvestment problem by making excellent output visible to leadership.
- Build first, pitch second: a working prototype beats a strategy memo every time.

- The "always-on" competitive monitoring model replaces quarterly review cycles with continuous intelligence.
- Start small—three competitors, one product line, one daily sales brief—and let demand pull you forward.

Positioning and Messaging in an Agent-Readable World

Pragmatic Remix: Positioning • Messaging • Buyer Personas • Use Cases • Buying Process

We want to tell you about two companies that sell roughly the same thing. Both offer enterprise data platforms. Both serve Fortune 500 customers. Both have strong engineering teams and credible analyst coverage. One of them—let's call them Company A—built their platform natively over fifteen years. The data fabric, the analytics engine, the planning tools, the AI layer—all designed to work together from the ground up, sharing a common semantic model and a unified governance framework. The other—Company B—assembled their platform through a series of acquisitions over five years. They bought a data warehouse company, a BI company, a data integration company, and an AI startup, then stitched them together with APIs and a unified brand and a marketing narrative about "one platform."

In the traditional buying process—the one where a human evaluation committee spends three months comparing vendors—both companies could compete effectively. Company B's sales team was polished. Their demo was carefully choreographed to hide the seams between acquired products. Their positioning emphasized breadth: "everything you need, one vendor, one contract." And it worked, because the human buyers were evaluating at the level of narrative and relationship and perceived risk reduction. They weren't parsing the technical architecture; they were responding to the story.

Then the agents showed up.

Here's what an agent does when a procurement team asks it to evaluate enterprise data platforms. It reads the documentation—all of it, not the curated demo path but the full technical docs, the API reference, the integration guides, the knowledge base articles. It notices that Company B's documentation references three different authentication models across three acquired products. It notices that the "unified semantic model" described on the marketing page doesn't match the technical reality described in the integration guides, where data moving between the acquired BI tool and the acquired data warehouse requires an ETL step that the marketing materials don't mention. It notices that the "one platform" positioning is contradicted by the support documentation, which has separate troubleshooting paths for each acquired component.

The agent doesn't care that Company B's sales team is charming. It doesn't care about the golf outing or the steak dinner or the executive sponsor who went to college with the prospect's CTO. It cares about structural coherence—whether the technical architecture matches the

marketing claims—because structural coherence is what it can evaluate at scale with high confidence. And on that dimension, Company A wins every time.

This is what we mean by "acquisition as confession." When a company acquires five products to build a platform, they're implicitly confessing that they didn't have the capability natively—and in an agent-mediated buying process, that confession is legible in the technical documentation even if it's invisible in the marketing narrative.

• • •

We're not telling this story to pick on any specific vendor—though if you've spent time in the enterprise data space, you can probably map Company A and Company B to real companies. We're telling it because it illustrates the most important thing that's changing about positioning in the agentic era: your positioning now has two audiences, and they evaluate it differently.

The Dual-Audience Problem

For as long as product marketing has existed, positioning has been a human-persuasion exercise. You write positioning to change how a person thinks about your product relative to alternatives. April Dunford's work—which is excellent, and which we'd recommend to any PMM who hasn't read *Obviously Awesome*—frames positioning as a deliberate act of contextualizing your product so that its value becomes obvious to the right buyer. That framing assumes a human buyer who processes information through narrative, analogy, social proof, and emotional resonance.

That assumption is no longer sufficient. Your positioning now needs to work for two audiences simultaneously.

Audience one is the human decision-maker—the CMO, the CTO, the VP of Data Engineering, whoever signs the contract. This audience still responds to narrative, still cares about brand, still wants to feel like they're making a smart choice that they can defend to their board. The human buyer hasn't disappeared. But they're increasingly making their decision based on a shortlist that was pre-filtered by audience two.

Audience two is the agent. The AI system that the buying organization has tasked with evaluating vendors, comparing capabilities, assessing technical fit, and producing a recommendation. This audience doesn't respond to narrative. It responds to structured information, consistent claims, verifiable evidence, and technical coherence. It doesn't care that you won a Gartner Leader position—unless the specific evaluation criteria in the Gartner report are relevant to the buyer's stated requirements, in which case it cares about the criteria, not the badge.

The dual-audience problem creates a tension that most PMMs haven't grappled with yet. Traditional positioning advice says to lead with the business outcome: "don't sell the drill, sell the hole." That's great advice for the human buyer. But the agent isn't evaluating whether you sell the hole effectively—it's evaluating whether your drill actually makes the hole. It wants specifications, integration capabilities, performance benchmarks, architectural details. It wants the receipt, not the promise.

What Agents Actually Read

We spent a couple of weeks running an experiment that we think every PMM should replicate. We took the positioning pages—the main product marketing pages, not the documentation—from ten enterprise software companies and ran them through Claude with a simple prompt: "You are an AI procurement agent evaluating data platform vendors for a mid-market manufacturing company. Based on this page, assess this vendor's fit for the following requirements." Then we listed five specific, technical requirements—data integration from ERP systems, real-time analytics, predictive planning capabilities, role-based access control, and multi-cloud deployment support.

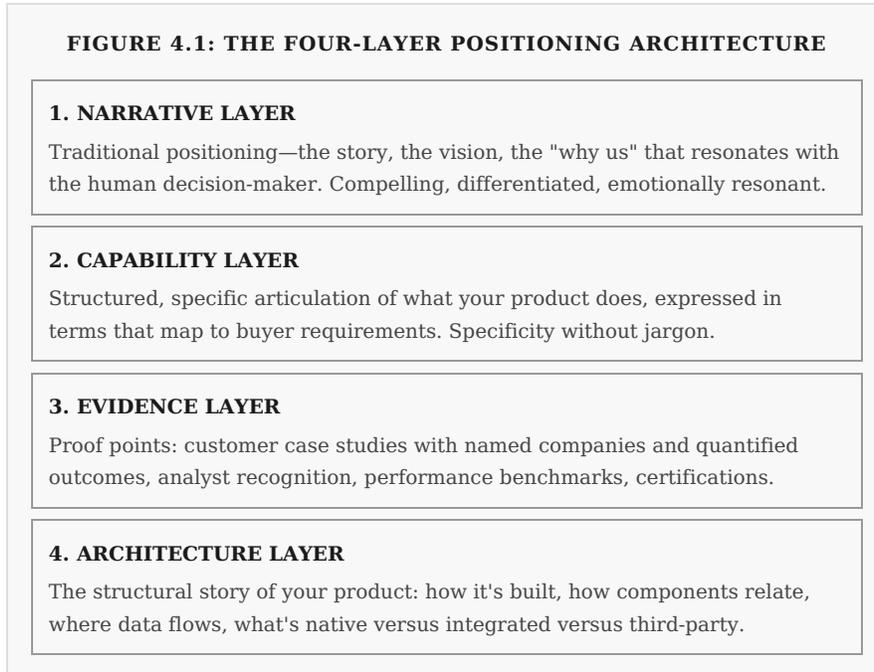
The results were revealing. About half of the product pages—the ones written in traditional marketing language, heavy on vision and light on specifics—produced agent evaluations full of qualifications: "The vendor appears to offer analytics capabilities but the page does not specify whether these are real-time." "Multi-cloud deployment is not explicitly confirmed; the language suggests cloud availability but does not enumerate supported platforms." The agent was essentially saying: *I can't evaluate this vendor because the marketing page doesn't give me enough structured information to work with.*

The other half—the pages that included specific capabilities, named integrations, architecture diagrams, and explicit statements about what the product does and doesn't do—produced clean, confident evaluations. The agent could map capabilities to requirements, identify gaps, and provide a clear recommendation. These pages weren't less compelling to humans; many of them were well-written and even narratively engaging. They just also happened to be specific.

The lesson isn't "write for robots." The lesson is that vague positioning—the kind that sounds great in a keynote but doesn't actually say what the product does—has always been a weakness, and agents are about to expose it at scale. Every PMM who has ever written "unlock actionable insights from your data" without specifying what kind of data, what kind of insights, and what "actionable" actually means in the context of the customer's workflow is about to discover that their positioning is invisible to the fastest-growing segment of their buyer base.

The Positioning Architecture

So what does good positioning look like in an agent-mediated world? We think it has four layers, and the trick is that each layer serves a different function and a different audience.



Layer one is the narrative layer. This is traditional positioning—the story, the vision, the "why us" that resonates with the human decision-maker. It's April Dunford territory. It should be compelling, differentiated, and emotionally resonant. It's what makes a CMO think "these are my people." Nothing about the agentic era diminishes the importance of this layer. If anything, it matters more, because the human decision-maker is going to encounter your narrative later in the buying process—after the agent has already shortlisted you—and the narrative is what tips the decision from "qualified vendor" to "preferred partner."

Layer two is the capability layer. This is the structured, specific articulation of what your product does, expressed in terms that map to buyer requirements. It's not a feature list—it's a capability framework that says: for this type of use case, our product provides this capability, integrated with these systems, supporting these workflows, at this scale. This is the layer that agents parse most effectively, and it's the layer that most marketing pages get wrong by being either too vague or too technical. The sweet spot is specificity without jargon: "Native bidirectional integration with SAP S/4HANA, Salesforce, and Workday for real-time data synchronization" is specific enough for an agent and clear enough for a human. "Seamless enterprise integration across your technology ecosystem" is neither.

Layer three is the evidence layer. This is where you substantiate your

positioning with proof points: customer case studies with named companies and quantified outcomes, analyst recognition with the specific evaluation criteria that drove it, performance benchmarks, security certifications, integration partnerships. Agents weight evidence heavily because it's verifiable—they can cross-reference your claimed G2 rating against the actual G2 page, they can check whether the analyst report you cite actually says what you claim it says. The evidence layer is where intellectual honesty pays off and where exaggeration gets caught.

Layer four is the architecture layer. This is the one that most PMMs don't think of as positioning, but it's increasingly the most important layer for technical buyers and the agents that serve them. It's the structural story of your product: how it's built, how the components relate to each other, where data flows, what's native versus integrated versus third-party. This is where the "acquisition as confession" dynamic plays out—an agent reading your architecture documentation can tell whether your platform is genuinely integrated or bolted together, and that assessment shapes everything downstream.

Killing the Vague Value Proposition

We have a confession of our own: we have written vague value propositions. Plenty of them. In the DMP era, we wrote things like "unlock the power of data-driven marketing" and "deliver the right message at the right time." Those phrases tested well in focus groups and looked great on a slide. They also said exactly nothing. They were marketing air—words arranged in a pattern that sounded meaningful but conveyed no information that could differentiate one vendor from another.

The agentic era is going to kill vague value propositions, and honestly, good riddance. When an agent is comparing five vendors on behalf of a buyer, it needs to evaluate each vendor against specific criteria. "Unlock actionable insights" doesn't map to any criteria. "Real-time anomaly detection across supply chain data with automated alerting and root-cause analysis" maps to several. The first is a wish. The second is a capability. Agents deal in capabilities.

This doesn't mean your positioning should read like a spec sheet. The narrative layer still matters—humans still buy from humans, and the agent is filtering, not deciding. But it does mean that every claim in your positioning needs to be backed by a specific capability, and every capability needs to be findable, parseable, and verifiable. If an agent can't find evidence that you actually do the thing you claim to do, you're invisible to that agent. And invisible to the agent increasingly means invisible to the buyer.

The Practitioner's Playbook: Positioning for Two Audiences

Here's what we'd tell any PMM working on positioning right now.

First, run the agent test on your own product pages. Take your main product marketing page and feed it to Claude or ChatGPT with a prompt that simulates an agent evaluation: "You are evaluating this vendor for [specific use case] with these requirements: [list five to seven concrete requirements]. Based only on this page, assess the vendor's fit." If the agent comes back with a confident, accurate assessment, your positioning is working. If it comes back hedging—"the vendor appears to offer X but this is not explicitly confirmed"—you have a specificity gap.

Second, audit your documentation for consistency. One of the most damaging things we see—and we've seen it at every company we've worked at, including our own—is a disconnect between what the marketing page says and what the technical documentation says. Marketing says "unified platform." The docs describe three separate configuration workflows. Marketing says "real-time analytics." The docs describe a batch-processing pipeline with a fifteen-minute refresh interval. Agents will find these inconsistencies, and they will interpret them as credibility problems. A PMM who owns the documentation experience—not just the marketing page—is going to outperform a PMM who doesn't.

Third, build your positioning from the capability layer up, not the narrative layer down. The traditional approach is to start with the story—the vision, the differentiation, the "why"—and then figure out how to support it with proof points. We're suggesting the reverse. Start with the specific capabilities your product actually delivers. Map those capabilities to buyer use cases. Quantify the outcomes where you can. Then write the narrative layer on top of that foundation. The positioning will be more specific, more defensible, and more effective for both audiences—because it's grounded in what's real rather than what sounds good.

THE CMO PERSPECTIVE

Positioning is the activity where the gap between what PMMs produce and what the business needs is widest—and the agentic era is making that gap more visible. Strong positioning should serve multiple contexts with the same core narrative: winning a deal, briefing an analyst, defending a pricing decision, informing a product roadmap conversation. If the positioning only works for one of those, it's a messaging document for a specific audience, not positioning.

The four-layer framework—narrative, capability, evidence, architecture—maps to how CMOs evaluate quality. When positioning has all four layers, the work has been done. When it's strong on narrative but thin on capability and evidence, someone did the fun part and skipped the hard part. One practice we've adopted: positioning reviews now include a documentation check. If we claim a capability on the marketing page, the documentation page must support it. In an agent-mediated buying process, documentation consistency may be the highest-leverage

positioning activity you can do.

KEY TAKEAWAYS

- **Positioning must be multi-context:** usable across deals, analyst briefings, pricing defense, and roadmap conversations.
- **The four-layer test**—narrative, capability, evidence, architecture—is a practical quality gate for positioning documents.
- **Agent-mediated buyers evaluate documentation,** not just marketing pages—consistency between the two is critical.
- **Run a documentation audit:** if your marketing claims aren't backed by product docs, agents will notice the gap before human buyers do.

Sales Enablement and the Death of the Static Battlecard

Pragmatic Remix: Sales Enablement • Sales Tools • Presentations • Collateral

We want to describe a scene that every product marketer has lived, and that we've lived more times than we can count. It's 4:47 PM on a Tuesday. A Slack message from a sales rep—let's call him Marcus—lands in the PMM channel: "Hey, I'm going up against [Competitor X] in a deal tomorrow morning. Our battlecard is from Q3. Anything new I should know?"

Marcus doesn't want to read a twenty-page competitive analysis. He wants three things: what has changed since the battlecard was published, what objections he's likely to face based on this specific prospect's profile, and what trap-setting questions he can ask to reframe the conversation in his favor. He needs those three things by 8 AM tomorrow. And the PMM who receives this message is already in back-to-back meetings until 6 PM, has a launch brief due on Wednesday, and hasn't looked at Competitor X's latest product announcement because they've been heads-down on a different priority.

So what actually happens? Usually one of three things. The PMM stays late and manually pulls together a quick brief from whatever they can find—which is heroic but unsustainable. Or the PMM points Marcus to the existing battlecard and tells him it's mostly still accurate—which might be true but doesn't give Marcus what he needs. Or Marcus goes into the meeting underprepared and relies on general product knowledge, which works fine if the competitor hasn't changed their pitch but fails badly if they have.

This scene—the real-time, deal-specific, competitive-intelligence-on-demand request that the PMM can't possibly fulfill at scale using manual methods—is the single best argument for why the static battlecard is already dead. Not because battlecards aren't useful. Because the format assumes a world where competitive intelligence is stable enough to be captured in a document and refreshed on a quarterly cadence. That world is gone.

• • •

The Battlecard Is an Artifact. The Intelligence Is a System.

We tested this proposition directly. For an article on futureofpmm.com, we gave the same competitive battlecard brief to Claude and ChatGPT and

compared the results. Same competitor, same product, same format requirements. The exercise wasn't really about which AI was better—though the differences were instructive—it was about demonstrating what becomes possible when you treat competitive enablement as a generation problem rather than a document management problem.

The traditional battlecard workflow goes something like this: the competitive PMM gathers intelligence (analyst reports, product announcements, win/loss feedback, sales anecdotes), synthesizes it into a structured document (usually a PDF or a Confluence page with sections for overview, strengths, weaknesses, landmines, objection handling, and trap-setting questions), publishes it to the sales enablement platform, and then updates it somewhere between quarterly and never.

The agent-powered workflow inverts this entirely. The intelligence-gathering runs continuously—the monitoring system from Chapter 3 feeds into it. The synthesis isn't a document; it's a query. Marcus doesn't open a battlecard. He asks a question: "What do I need to know about Competitor X for a deal with a mid-market manufacturing company evaluating data platforms for supply chain analytics?" And the system generates a response that incorporates the latest competitive intelligence, tailored to the specific deal context, grounded in the same knowledge base that the static battlecard drew from but updated continuously and customized to the ask.

This isn't science fiction. We've seen it working. Soo Shim's BattleCoach at SAP does exactly this—and not in a demo environment, but in production, supporting live sales cycles. The system ingests the same intelligence sources that would have fed a traditional battlecard and makes them available as a conversational interface. Sales reps can ask questions in natural language. The system provides answers grounded in current competitive data. And because it's a system rather than a document, it doesn't decay—it gets better as more intelligence flows in.

The Enablement Paradox

There's an uncomfortable truth lurking beneath the sales enablement transformation, and we want to name it directly: sales reps are getting their own AI tools, and those tools are changing what "enablement" means.

Two years ago, when a sales rep needed to understand a competitor, they went to the PMM. When they needed to customize a pitch deck, they asked the PMM. When they needed a customer-facing one-pager tailored to a specific industry, the PMM created it. The PMM was the bottleneck—the human through whom all competitive and positioning intelligence flowed on its way to the customer.

Today, a sales rep with a Claude or ChatGPT subscription can generate a reasonably good competitive summary in thirty seconds. They can customize a pitch deck by pasting in the prospect's website and asking the

AI to tailor the messaging. They can produce an industry-specific one-pager that's—let's be honest—about 70% as good as what the PMM would have created. Maybe 80%.

This creates what we think of as the *enablement paradox*: the better AI tools get at producing competent sales content, the less sales teams rely on PMMs for that content—but the more they need PMMs for the thing that AI can't provide, which is the strategic intelligence and judgment that makes the content accurate and differentiated rather than merely competent.

A sales rep's AI can generate a competitive comparison. It can't tell you that the competitor's new partnership announcement is actually a defensive move because they're losing deals in the mid-market segment, and that the right response isn't to counter their messaging but to lean into the very use cases they're implicitly conceding by pivoting their strategy. That's PMM intelligence—the kind that comes from deeply understanding the competitive landscape, the buyer psychology, and the market dynamics. And in the agentic era, that intelligence is more valuable than ever because it's the input that makes all the other AI-generated content correct.

From Artifact Factory to Intelligence Layer

The role shift this implies is significant, and we want to be clear about it because we think a lot of PMMs are going to be surprised by how fast it happens.

In the old model, the PMM's sales enablement job was primarily artifact production. You created battlecards, pitch decks, one-pagers, talk tracks, objection-handling guides, customer reference lists, and demo scripts. The quality of those artifacts was how your sales enablement effectiveness was measured—and how, let's be real, you justified your headcount to leadership.

In the new model, the PMM's sales enablement job is primarily intelligence architecture. You design the systems that ensure sales teams have access to accurate, current, contextually relevant competitive and product intelligence at the moment they need it. You curate the knowledge base that the AI tools draw from—because the quality of the AI output is only as good as the intelligence that feeds it. You build the feedback loops that surface what's working in deals and what isn't. You're the person who ensures that when Marcus asks his AI for help with Competitor X, the answer is based on yesterday's intelligence, not last quarter's.

The artifacts still exist. They're just generated, not crafted. And the PMM's job shifts from producing the artifacts to ensuring the intelligence substrate that feeds them is accurate, current, strategically grounded, and differentiated. That's a harder job than producing battlecards. It's also a more valuable one.

The Knowledge Base as Competitive Moat

Here's something that we don't think enough PMMs have internalized yet: the quality of your internal knowledge base is about to become one of your most important competitive advantages.

In a world where both your sales reps and your competitors' sales reps have access to the same general-purpose AI tools, the differentiator isn't the tool—it's the data the tool draws from. If your competitive knowledge base is deep, current, well-organized, and rich with contextual intelligence (win/loss insights, customer verbatims, deal-specific lessons, analyst perspectives), your AI-powered enablement system produces output that is categorically better than what your competitor's rep gets from the same general-purpose AI.

This is where the RAG architecture—retrieval-augmented generation—becomes a PMM responsibility, not just a technical architecture. RAG is the pattern where an AI retrieves relevant information from a curated knowledge base before generating a response, grounding its output in specific, verified data rather than general training knowledge. We wrote a weekend-project tutorial on futureofpmm.com about building a personal RAG system, and the response surprised us. PMMs at companies from Series B startups to Fortune 100 enterprises reached out to say they were building something similar for their sales enablement workflows.

The pattern is straightforward. You build a knowledge base—this can be as simple as a well-organized collection of documents in a vector database, or as sophisticated as a purpose-built system with metadata tagging, recency weighting, and source quality scoring. You populate it with your competitive intelligence, your positioning documents, your win/loss analyses, your customer case studies, your product documentation, and your analyst coverage. You connect it to a generation layer that can synthesize this knowledge into deal-specific outputs on demand. And then you maintain it—because a RAG system is only as good as the knowledge base behind it, and an out-of-date knowledge base produces confidently wrong answers, which is worse than no answer at all.

The PMM who owns and curates this knowledge base—who makes it their job to ensure that the intelligence substrate is excellent—is the PMM who becomes indispensable. Not because they're personally producing every battlecard and one-pager, but because every AI-generated enablement asset in the organization draws from the intelligence they maintain.

The Practitioner's Playbook: Rebuilding Sales Enablement

If you're a PMM responsible for sales enablement, here's where we'd start.

First, audit your knowledge base—and be honest about what you find. Pull up your competitive battlecards, your positioning documents,

your product one-pagers, your customer references. When was each one last updated? How many of them are still accurate? If you fed them to an AI system right now, would the AI produce good output, or would it produce confidently stale output? In our experience, most PMM teams discover that 40-60% of their enablement content is materially out of date. That's not a failure of effort; it's a failure of model. The quarterly-refresh model can't keep pace with a market that shifts weekly.

Second, pick one competitor and build the full pipeline. Don't try to boil the ocean. Choose your most important competitor—the one that shows up in the most deals—and build an end-to-end agent-powered enablement workflow for that competitor alone. Set up continuous monitoring. Build a RAG-connected knowledge base with everything you know about them. Create a natural-language interface—even if it's just a well-engineered Claude prompt with your documents loaded as context—that sales reps can query. Get three or four reps to test it for two weeks and give you feedback. Iterate. Then expand to the next competitor.

Third, redesign your enablement metrics. If your current metrics are artifact-based—number of battlecards produced, content downloads, deck views—they're measuring the wrong thing. In the agent-powered model, the metrics that matter are: how current is the intelligence (average age of content in the knowledge base), how frequently is it being used (queries per week from the sales team), and most importantly, what's the outcome (win rate in competitive deals where the agent-powered enablement was used versus deals where it wasn't). That last metric is hard to measure cleanly, but even a directional signal is more useful than counting PDF downloads.

THE CMO PERSPECTIVE

The enablement paradox is real, and it's happening faster than most PMM leaders expect. Sales teams have adopted AI tools almost overnight—in some cases faster than the PMM teams supporting them. Reps are using ChatGPT to prep for calls, generate email sequences, and produce quick competitive summaries. The natural PMM response is concern: "If sales can do this themselves, what do they need us for?" The answer is intelligence—the strategic layer that determines whether AI-generated content is right or wrong, differentiated or generic, helpful or dangerous.

From a budget and headcount perspective, this shift is actually clarifying. Justifying PMM value based on artifact output—twelve battlecards per quarter, four product launches—always confused activity with impact. In the agent-powered model, enablement effectiveness becomes more directly measurable: Is sales using the intelligence system? Are win rates improving in competitive deals? Is the time from competitive signal to sales response shrinking? Those are metrics that tie directly to revenue, and that's a better position for the PMM function, not a worse one.

KEY TAKEAWAYS

- **Sales reps are already using their own AI tools**—the PMM's job shifts from producing content to ensuring the intelligence layer is accurate.
- **The knowledge base, not the battlecard**, becomes the primary enablement artifact.
- **New PMM metrics:** intelligence system adoption, competitive win rate improvement, signal-to-response time.
- **The enablement paradox resolves in the PMM's favor** when value is measured by impact rather than artifact volume.

Content, Thought Leadership, and the Narrative Advantage

Pragmatic Remix: Thought Leadership • Content Creation • Customer Stories • Collateral

In early 2026, we did something slightly masochistic: we read every article about AI and product marketing that appeared in a Daily Digest feed for an entire month. Thirty-five editions. A hundred and eighty articles. LinkedIn posts, analyst reports, vendor blogs, trade press, Substack newsletters, conference recaps. We read them all, and then wrote about what we found on futureofpmm.com. The piece was called "I Read Every AI + PMM Story for a Month. Here's What's Actually Coming."

The experience was clarifying in a way we didn't expect. Not because the articles were bad—some were excellent—but because the pattern was so stark. About 80% of the content fell into a category we'd describe as *competent and interchangeable*. Solid writing, reasonable arguments, accurate information, and absolutely nothing that would stick in your memory twenty minutes after reading it. These were articles that could have been written by any of a hundred different people—or, increasingly, by no person at all. They hit all the standard beats: AI is transforming marketing, PMMs need to adapt, here are five things to think about, here's a framework. The frameworks were always reasonable. The advice was always sensible. And the total information content, after you'd read the third or fourth article making the same argument with different metaphors, was approximately zero.

The other 20% was different. These were pieces where the author had a specific experience, a contrarian take, a story nobody else had told, or a level of technical depth that couldn't be faked. A PMM at AppDirect describing how her team built eleven SaaS applications and twenty AI agents in a year—not because she was told to but because she made a ridiculous promise to her team and then figured out how to keep it. An analyst at Gartner threading together three different market trends into a prediction that contradicted the consensus. A founder writing about why their company was deliberately not building an AI feature that every competitor was rushing to ship.

The ratio—80/20—is about to get worse. Much worse. Because the 80% that's competent and interchangeable is exactly the content that AI produces most easily. And the marginal cost of producing it is dropping toward zero.

The Content Collapse

Let us put some numbers around this, because the scale of what's happening is hard to grasp intuitively.

In 2023, a competent PMM might produce two to three substantive blog posts per month, a handful of social posts per week, and a quarterly thought leadership piece. That's maybe 15,000 words of original content per month, and it represented a significant chunk of their working time. In 2026, that same PMM—armed with Claude or ChatGPT and a decent prompting workflow—can produce that volume in a day. Not at the same quality, maybe. But at a quality level that's good enough for most distribution channels and indistinguishable from human-written content to most readers.

Now multiply that by every PMM at every company in your competitive landscape. The supply of competent marketing content is exploding while the demand—the finite attention of your buyers—stays the same. That's the *content collapse*: the moment when the sheer volume of AI-generated, competent-but-generic content overwhelms every distribution channel and makes it dramatically harder for any individual piece to break through.

We've heard people describe this as "thought leadership is dead," and the phrase has been popping up in the kind of LinkedIn posts that are themselves examples of the problem. But the diagnosis is only half right. What's dying is *performative* thought leadership—content that performs the appearance of insight without delivering any. The listicles, the framework-of-the-week posts, the "5 Things Every CMO Needs to Know About AI" articles that all contain the same five things. That genre is dying because AI can produce it at scale and because readers are developing antibodies to it.

What's not dying—what's becoming scarcer and therefore more valuable—is *genuine* thought leadership. Content where the author has done something, seen something, learned something, or thought something that readers can't get anywhere else. Content with a specific point of view, grounded in specific experience, delivered with a voice that is recognizably human. That kind of content was always rare. Now it's the only kind that matters.

The Narrative Advantage

We've written three books about data-driven marketing, and every one of them was ultimately about the same problem: how do you take something that sounds boring on paper—data management platforms, identity resolution, customer data architectures—and make people care about it? The answer was always stories. Not stories as decoration, pasted onto a technical argument to make it go down easier. Stories as the argument itself.

In *Data Driven*, the story that made people care about DMP economics was Mike Cunningham at Keurig wanting to put a \$5 chip in every coffee maker to understand household-level consumption patterns. Dad drinks dark roast Sumatra at 6 AM, Mom prefers French Hazelnut at 7:30, and the twelve-year-old sneaks Donut Shop K-Cups on weekends. You remember that image. You might not remember the TAM analysis from the same chapter, but you remember the family and the coffee and the chip. In *Customer Data Platforms*, the story that made identity resolution concrete was the Porsche sitting on the tarmac at the Atlanta airport—a Delta SkyMiles loyalty program data play where they knew exactly which high-value customers to target with a once-in-a-lifetime upgrade experience. The technical explanation of deterministic versus probabilistic matching would have put you to sleep. The Porsche on the runway woke you up.

AI can't do this. Not yet, and maybe not ever—or at least not in the way that matters. An LLM can generate a competent analogy. It can produce a hypothetical scenario that illustrates a point. But it can't tell you about Douwe Bergsma interrupting a deck to ask why a toilet paper company needs a DMP, because it wasn't in the room. It can't tell you how Mike Cunningham's eyes lit up when he talked about the Keurig chip, because it didn't see it happen. The specificity of lived experience—the named person, the exact question, the energy in the room—is the thing that makes a story stick, and it's the thing that AI-generated content systematically lacks.

This is what we call the *narrative advantage*, and we believe it's going to be one of the most important differentiators for PMMs in the agentic era. The PMM who has stories—real stories, from real customer conversations, real deal cycles, real product failures and recoveries—and who can tell those stories in a way that makes abstract capabilities feel concrete and human, has something that no AI agent can replicate. Not because the technology isn't good enough. Because the stories require having been somewhere and paid attention to what happened there.

GEO and the Agent as Reader

There's another dimension to the content transformation that most PMMs aren't thinking about yet, and it's going to matter as much as the thought leadership question. It's called *generative engine optimization*—GEO—and it's the discipline of optimizing your content to be surfaced and cited by AI systems rather than (or in addition to) traditional search engines.

We profiled a company called Evertune on futureofpmm.com—they're building what we described as the Bloomberg Terminal for AI search. The founders came from The Trade Desk, which means they understand how platforms get built and how measurement changes behavior. Their thesis is that GEO will become one of the most important channels for product marketing, and that most companies are completely blind to how they're showing up—or not showing up—in AI-generated answers.

Think about what this means for your content strategy. When a buyer asks ChatGPT or Claude or Perplexity to recommend data platform vendors for a specific use case, the AI synthesizes its response from the content it can access—web pages, documentation, analyst reports, review sites, blog posts. If your content is specific, structured, and authoritative on the topics that matter to that buyer, you get cited. If your content is generic or vague or buried behind a registration wall, you don't. In a world where a growing percentage of vendor discovery happens through AI-mediated search rather than traditional Google search, your content isn't just competing for human attention—it's competing for inclusion in the AI's synthesis.

This connects directly to the positioning argument from Chapter 4. The same principles that make your positioning agent-readable—specificity, structured claims, verifiable evidence—also make your content more likely to be surfaced by AI systems. The PMM who writes a blog post titled "Unlocking the Power of Data-Driven Insights" is invisible to GEO because the content says nothing specific enough to be cited in response to a concrete buyer question. The PMM who writes "How We Reduced Supply Chain Forecasting Errors by 34% Using Real-Time ERP Data" is exactly what an AI system surfaces when a buyer asks about supply chain analytics capabilities.

The Content Operating Model

So what does a content strategy actually look like in this environment? We think it has three layers, and the time allocation between them is roughly the inverse of what most PMM teams practice today.

The first layer is commodity content—the blog posts, social updates, email sequences, slide decks, and one-pagers that keep the marketing engine running. This layer should be almost entirely agent-produced. Not because it doesn't matter, but because the marginal return on human effort here is approaching zero. An agent can produce a LinkedIn post about your product launch that's 90% as good as what a human would write, in 2% of the time. The human's job is to review it, ensure it's accurate, and add a line of personality if the occasion warrants it. But the drafting, the formatting, the adaptation across channels—all of that is agent work now. If your PMM team is spending more than 20% of their content time on commodity production, they're misallocating.

The second layer is strategic content—the pieces that advance your positioning, establish your point of view, and create the intelligent context that agents and buyers draw from when evaluating your product. This is where the GEO investment lives: deeply specific articles about capabilities, use cases, architectural approaches, and customer outcomes. It's also where competitive content lives—the pieces that articulate why your approach is different and why it matters. This layer is human-directed and agent-assisted. The PMM defines the argument, provides the specific

stories and data points, and ensures the strategic positioning is correct. The agent handles research, first-draft generation, and format adaptation. We'd allocate about 40% of content time here.

The third layer is signature content—the pieces that can only come from a specific person with specific experiences and a specific point of view. These are the conference keynotes, the long-form essays, the genuinely original frameworks, the customer stories told with enough texture and detail that you can feel the room. This layer is entirely human. No agent assistance in the writing—though agents might help with research or fact-checking—because the point of signature content is that it's irreducibly personal. It's what makes readers follow a specific author rather than a brand. It's what makes analysts quote your perspective rather than your company's positioning. It's the 20% that disproportionately drives the other 80%.

If you invert those percentages from where most PMM teams are today—where 60% of time goes to commodity content, 30% to strategic content, and 10% to signature content—and move toward a 20/40/40 split, you've transformed the content operation from a production function into a strategic one. That's the 10x move for content.

The Practitioner's Playbook: Content in the Agentic Era

Three things to do in the next thirty days.

First, run the GEO audit. Ask Claude, ChatGPT, and Perplexity to recommend products in your category for three to five specific use cases that your buyers care about. Do you show up? Are you cited? Is the information accurate? If you're invisible or misrepresented, that's a content gap with direct pipeline implications. The fix is almost always more specific, more structured, more technically detailed content that addresses the exact questions buyers ask.

Second, identify your signature content and protect the time for it. Look at the content your team has produced in the last year and find the pieces that got the most engagement—not just views, but comments, shares, inbound from prospects, mentions by analysts. We'll bet those pieces are the ones with the most specific stories, the strongest points of view, and the most recognizable voice. That's your signature content. Now look at how much time the person who wrote those pieces has to produce more of them. If the answer is "not enough because they're too busy producing commodity content," you've identified the problem.

Third, build the commodity content pipeline. Choose your LLM, define your templates, create a prompting workflow that takes your core messaging and produces channel-specific content at the quality level you need. This isn't a one-time project; it's an operating system. The goal is to reduce commodity content production to a review-and-approve workflow rather than a draft-from-scratch workflow. The hours you reclaim go

straight into strategic and signature content.

THE CMO PERSPECTIVE

The three-layer content model—commodity, strategic, signature—matches how content investment decisions should work at the CMO level. The conversation with leadership about content has changed. It used to be about volume: how many blog posts per month, how many assets per launch. Those were easy metrics to track and easy to game. A team could hit volume targets and still produce nothing that moved the needle on pipeline or perception.

The shift is toward influence metrics. Did the analyst cite our framework in their report? Did the prospect reference our content in the sales call? Did the competitive positioning piece change how sales talks about us in deals? These are harder to measure but they're the metrics that indicate strategic impact. In a fifty-plus-person PMM organization, maybe five people produce genuinely differentiated signature content—content where the author's experience and voice are integral to the argument. Those five are disproportionately valuable, and agent-powered commodity production is the biggest lever for freeing them up.

KEY TAKEAWAYS

- **Content measurement should shift from volume metrics** (posts per month) to influence metrics (analyst citations, sales deal impact).
- **The commodity/strategic/signature framework** is also a resource allocation model: automate the bottom, invest at the top.
- **Signature content**—where human voice and experience are integral—is the scarcest and highest-return content investment.
- **Agent-powered commodity production** frees your best thinkers to do more of the strategic and signature work that moves markets.

The Demo Stack: Product Storytelling Reimagined

Pragmatic Remix: Product Demos • Use Case Development • Proof of Concept

Have you ever done a product demo? We don't mean watched one—we mean delivered one. Stood in front of a room, or more likely a Zoom screen, and tried to make enterprise software feel exciting? If you have, you know the particular agony of the moment when the demo environment crashes, or the data doesn't load, or the feature you're about to show requires a login that expired overnight. You're mid-sentence, the prospect is leaning forward, and you're silently praying to whatever deity governs staging environments.

Demos have always been product marketing's highest-stakes deliverable. A great demo closes deals. A bad demo—or worse, a broken demo—can kill months of pipeline development in ninety seconds. And yet the infrastructure around demos has historically been shockingly primitive. Most enterprise companies run demos on shared staging environments that are one bad deployment away from embarrassment. The "demo script" is usually a Google Doc that gets updated when someone remembers to update it. The personalization—tailoring the demo to the prospect's industry, use case, and technical environment—is done manually by the sales engineer or the PMM, often at 11 PM the night before.

This is the problem that a new category of tools—the demo automation stack—is solving. And the agentic era is about to accelerate the transformation from demos-as-live-performance to demos-as-interactive-experience in a way that changes the PMM's role entirely.

The Demo Landscape

We did a comprehensive evaluation of demo automation platforms for futureofpmm.com—the piece was called "The Demo Stack"—and the landscape is more mature and more varied than most PMMs realize. The tools cluster into a few categories, and understanding the categories matters because they serve different stages of the buyer journey and require different levels of PMM involvement.

The first category is demo environment platforms—tools like Demostack and Reprise that let you create sandboxed, customizable versions of your actual product. These aren't screenshots or click-through prototypes; they're functional replicas of your application with realistic data that you control. The sales engineer or PMM can customize the demo environment for a specific prospect—swap in their industry's data, configure the

workflows they care about, pre-load the integrations they use—without touching the production environment. The demo works reliably because it's decoupled from the development team's release cycle.

The second category is interactive demo builders—tools like Navattic, Supademo, and Arcade that create guided, self-serve product tours. These are lighter-weight than full environment replicas but more sophisticated than static screenshots. They let you build a clickable walkthrough of key workflows with annotations, tooltips, and branching paths. The prospect can explore the product at their own pace, on their own schedule, without waiting for a sales call. For top-of-funnel engagement—where the buyer is evaluating whether your product is worth a deeper look—these are increasingly replacing the live demo as the first product experience.

The third category is demo overlay tools—platforms like Saleo and Walnut that let you overlay customized data and narratives onto your actual product interface. Instead of building a separate demo environment, you run the demo in your real product but with a layer on top that controls what data appears, how it's labeled, and what story it tells. This is particularly useful for complex enterprise products where the full demo environment approach would require too much setup.

The fourth—and this is the one that intersects most directly with the agentic thesis—is *asynchronous demo platforms* like Consensus, which let you create personalized video demos that buyers can watch and interact with on their own. These platforms track engagement at the individual level: which features did the buyer spend time on, which sections did they skip, who else at the buying organization watched the demo. That engagement data feeds back into the sales process as intent signals.

What Is the Demo For?

Here's the question that the agentic era forces us to ask: if an AI agent can evaluate your product's technical capabilities from your documentation, your API specs, and your architecture diagrams—and increasingly, it can—what is the demo actually for?

In the traditional model, the demo served two functions: *proof of capability* (does the product actually do what the marketing claims?) and *proof of experience* (what does it feel like to use this product?). The first function is being absorbed by agents. When a buyer's AI agent can read your API documentation, test your integration endpoints, and verify that the capabilities you claim actually exist, the demo's role as capability proof diminishes. The prospect's agent already knows whether your product supports real-time data synchronization with SAP S/4HANA before the sales call starts.

The second function—proof of experience—becomes the demo's primary job. And that's a shift that plays directly to the PMM's strengths, because proof of experience is a storytelling problem, not a technical problem.

The best demo we've ever watched—and we've watched hundreds—wasn't a feature walkthrough. It was a narrative. The sales engineer started with the customer's current state: "Right now, your supply chain team gets a forecast update every Monday morning. They download a spreadsheet, manually reconcile it against the ERP data, and by Wednesday they have a planning view that's already two days stale." Then she showed the same workflow in the product—not as a feature tour but as a transformation story. "Here's what Monday morning looks like with our platform. The forecast is live. The reconciliation is automatic. And your supply chain team spends Wednesday on strategy instead of spreadsheets." The prospect didn't just see what the product did. They felt what their life would be like using it.

That's the demo that matters in the agentic era—the one that creates an emotional and experiential connection that no agent evaluation can replicate. The agent can tell you whether the product meets the technical requirements. Only the demo can make you want it.

Personalizing Demos at Scale

The obvious challenge with the "demo as narrative" approach is that it requires personalization, and personalization doesn't scale in a traditional workflow. Building a customized demo narrative for every prospect—their industry, their specific pain points, their current tech stack, their organizational context—is the kind of work that would consume a PMM's entire week if they tried to do it for more than a handful of deals.

This is where agent augmentation makes the most difference. The same intelligence infrastructure that powers your competitive enablement (Chapter 5) can power your demo personalization. An agent can pull the prospect's company data—industry, size, tech stack, recent news—and combine it with your product's use case library to generate a customized demo script that tells the right story for that specific buyer. The PMM or sales engineer reviews and adjusts the script—they add the specific anecdote that makes it feel human, they adjust the emphasis based on what they know about the buying committee's priorities—but the heavy lifting of research and first-draft generation is handled by the agent.

Paired with the demo automation platforms we described earlier, this creates a workflow where a personalized, industry-specific, narrative-driven demo can be prepared in hours instead of days. The demo environment is pre-configured with relevant data. The script is drafted around the prospect's specific use case. The sales engineer's job shifts from building the demo to performing it—and performing it well, which is the part that actually wins deals.

The Practitioner's Playbook: Rethinking Your Demo Strategy

If you own the demo experience for your product, here's what we'd prioritize.

First, separate the proof of capability from the proof of experience.

Make sure your documentation, your API specs, and your product pages are detailed enough that an agent can evaluate your technical capabilities without needing a live demo. This isn't just good for agent-mediated buying; it also filters out the tire-kickers from your demo pipeline. If a prospect's agent can determine that your product doesn't meet a key requirement before the demo is scheduled, you've saved everyone's time.

Second, invest in a demo automation platform if you haven't already.

The category has matured enough that there are good options at every price point and complexity level. For self-serve, top-of-funnel experiences, interactive tour builders like Navattic and Supademo are the right starting point—they're relatively inexpensive and can be built by a PMM without engineering support. For high-value enterprise demos, environment platforms like Demostack and Reprise are worth the investment because they eliminate the staging-environment-crash risk and enable the kind of deep personalization that wins deals.

Third, build a demo narrative library.

Document the five to seven transformation stories that resonate most with your buyers—the before/after narratives that connect your product's capabilities to the prospect's lived experience. These aren't feature lists; they're stories with characters, tension, and resolution. Once you have the library, an agent can customize each narrative for a specific prospect, and the sales engineer has a foundation to build on rather than starting from scratch every time.

THE CMO PERSPECTIVE

The demo investment question is one of the harder ROI cases to make—not because the return isn't there but because the metrics are indirect. A better demo experience shows up as a shorter sales cycle, a higher conversion rate from qualified pipeline to closed-won, and anecdotal feedback from sales about prospects who "just got it" faster. What's easier to quantify is the cost of not investing: deals lost or delayed because the demo environment broke, hours sales engineers spend on prep versus selling, prospects who requested self-serve product experiences that didn't exist.

The demo as narrative experience—rather than capability proof—is the key unlock. The best sales engineers already do this instinctively: they tell the customer's transformation story, not the feature story. Demo automation platforms make the narrative approach scalable by separating the story from the environment, so even a newer SE can deliver a compelling, personalized demo built on the same framework the top performers use.

KEY TAKEAWAYS

- **Demo ROI is best measured by absence:** lost deals, wasted SE hours, and unmet self-serve requests quantify the cost of underinvestment.
- **The demo should be a narrative experience**—the customer's transformation story, not a feature walkthrough.
- **Demo automation separates story from environment,** making top-performer quality scalable across the sales team.
- **Evaluate demo tools honestly:** avatar and video tools are not the same as true interactive demo platforms.

Pricing, Packaging, and the Consumption Shift

Pragmatic Remix: Pricing • Packaging • Business Case • Profitability

We're going to start this chapter with a question that sounds simple and isn't: when an AI agent performs a task on behalf of a user, who is the user?

This isn't a philosophical question. It's a pricing question—and it's one that every enterprise software company is wrestling with right now, whether they've articulated it or not. The per-seat licensing model that has been the foundation of SaaS economics for two decades is built on a simple assumption: value is proportional to the number of humans who use the product. More users equals more value equals more revenue. The entire pricing architecture—the sales motion, the expansion playbook, the financial model, the investor narrative—flows from that assumption.

Agents break the assumption. When a single user can deploy an AI agent that performs work equivalent to what five users used to do, the per-seat model collapses. The user count goes down while the value delivered goes up. Under per-seat pricing, this is a revenue problem. Your product is delivering more value and getting paid less for it.

We've spent a significant part of our time at SAP working on this exact problem—figuring out how to transition pricing and packaging models for data and analytics products from traditional SaaS to consumption-based and AI-credit frameworks. It's one of the hardest problems in product marketing, and it's one where most PMMs have almost no training, because pricing has historically been treated as a finance-and-product decision rather than a marketing decision. That's a mistake. *Pricing is positioning*. How you charge for something tells the market what you think it's worth, and how it should be valued relative to alternatives.

The Three Pricing Models

In the agentic era, enterprise software pricing is converging on three models. Most companies will end up with some combination of all three, but understanding each in isolation is useful for thinking about where the market is heading.

This isn't a philosophical question. It's a pricing question — and it's one that every enterprise software company is wrestling with right now, whether they've articulated it or not. The per-seat licensing model that has been the foundation of SaaS economics for two decades is built on a simple assumption: value is proportional to the number of humans who use the product.

Pricing is positioning. How you charge for something tells the market what you think it's worth, and how it should be valued relative to alternatives.

Figure 1: The Three Pricing Models

Three models are converging in enterprise software. Most companies will deploy a hybrid — but understanding each in isolation clarifies the strategic trade-offs.

Dimension	Per-Seat	Consumption	Outcome / AI-Credit
Value alignment	Weak. More users ≠ more value when agents do the work.	Strong. Pay scales with actual usage of compute, storage, queries.	Strongest. Customer pays for results delivered, not resources consumed.
Predictability	High. Fixed cost per user per month. CFOs love it.	Medium. Requires committed-spend tiers or capacity blocks to smooth.	Low. Defining and measuring "outcomes" is genuinely hard.
Complexity	Low. Simple to sell, simple to buy, simple to model.	Medium. Requires usage metering, alerts, and spend governance.	High. Requires outcome definition, measurement, and attribution.
Best for	Collaboration tools, comms platforms — where AI hasn't changed work-per-user ratio.	Data platforms, analytics, infrastructure — usage-proportional value.	AI-native products where the agent does the work and the human reviews.
Competitive risk	High. Competitors moving to consumption expose your model as misaligned.	Medium. Must package well or customers fear unpredictable bills.	Low risk, high execution bar. First mover advantage if you get it right.
PMM challenge	Defend the model or plan the transition.	Design packaging that balances usage alignment with spend predictability.	Define what an "outcome" is, price it, and make it sellable.

Figure 1. The three pricing models of the agentic era. Per-seat is stable but increasingly misaligned. Consumption is proven but requires packaging discipline. Outcome-based is the frontier — philosophically right, operationally hard.

The second is consumption-based pricing, where the customer pays for what they use — typically measured in compute, storage, queries, API calls, or some other usage metric. This model aligns revenue with value delivered: the more a customer uses the product, the more they pay. Snowflake popularized this model in the data platform space, and it's spreading across enterprise software.

Figure 1: The three pricing models of the agentic era.

The first is the traditional per-seat model, which still works for products where the value is proportional to the number of human users and where AI augmentation hasn't fundamentally changed the work-per-user ratio. Collaboration tools, communication platforms, and certain workflow

products still fit this model—though even there, the pressure is building as AI features reduce the number of humans needed to perform a given task.

The second is consumption-based pricing, where the customer pays for what they use—typically measured in compute, storage, queries, API calls, or some other usage metric. This model aligns revenue with value delivered: the more a customer uses the product, the more they pay. Snowflake popularized this model in the data platform space, and it's spreading across enterprise software. The challenge with consumption pricing is predictability—customers want to forecast their spend, and a purely consumption-based model makes that difficult. The PMM's job is to help design packaging that provides enough predictability (committed spend tiers, volume discounts, pre-purchased capacity blocks) without eliminating the alignment between usage and revenue.

The third—and this is the frontier—is outcome-based or AI-credit pricing, where the customer pays for the results that AI agents deliver rather than the resources consumed. An AI credit might represent a specific unit of agent work: a competitive analysis produced, a forecast generated, a document processed, a decision recommended. This model is the most philosophically aligned with the agentic era—the customer pays for value delivered, not for human seats or machine resources—but it's also the hardest to implement because defining and measuring "outcomes" is genuinely difficult.

Why Pricing Is a PMM Problem

In most organizations we've worked in, pricing decisions are made by a triangle of product, finance, and executive leadership, with product marketing providing "inputs"—competitive pricing analysis, customer willingness-to-pay research, packaging recommendations—but not sitting at the decision table. That model is inadequate for the pricing complexity of the agentic era, and PMMs need to fight for a bigger seat.

Here's why. Pricing in the agentic era isn't just a financial engineering problem. It's a positioning problem, a competitive strategy problem, and a buyer psychology problem—all of which are core PMM competencies. When you choose between per-seat and consumption pricing, you're making a positioning statement about who your product is for and how it delivers value. When you design packaging tiers, you're making a competitive strategy statement about which market segments you're targeting and which ones you're willing to concede. When you communicate the pricing model to buyers, you're navigating buyer psychology around risk, predictability, and perceived fairness.

Pricing in the agentic era isn't just a financial engineering problem. It's a positioning problem, a competitive strategy problem, and a buyer psychology problem — all of which are core PMM competencies. When you choose between per-seat and consumption pricing, you're making a positioning statement about who your product is for and how it delivers value.

When you choose between per-seat and consumption pricing, you're making a positioning statement about who your product is for and how it delivers value. When you design packaging tiers, you're making a competitive strategy statement about which market segments you're targeting and which ones you're willing to concede.

Figure 2: The Pricing-Positioning Connection

Every pricing decision is three decisions at once. The PMM who treats pricing as "just a finance thing" is ceding positioning, competitive strategy, and buyer psychology to people who don't own those disciplines.

PRICING MODEL	POSITIONING SIGNAL	COMPETITIVE IMPLICATION	BUYER PSYCHOLOGY
Per-Seat	"Our product is a tool for people. Value comes from human users doing human work."	Vulnerable to competitors who price on value delivered. Defenders must articulate why human-in-the-loop matters.	Feels safe and predictable. Buyers know exactly what they'll spend. But increasingly feels misaligned — "Why am I paying for seats my agents are using?"
Consumption	"Our product delivers value proportional to how much you use it. We're confident you'll use it a lot."	Advantage against per-seat competitors in high-usage scenarios. Requires packaging to counter "unpredictable spend" objection.	Feels fair but risky. Buyers want transparency on what drives consumption and guardrails against runaway spend.
Outcome / AI-Credit	"Our product delivers measurable results. We're so confident in the value that we'll tie our revenue to your outcomes."	Strong differentiation if credible. Forces competitors to either match or explain why they won't tie pricing to results.	Feels aspirational — buyers want this. But skepticism is high: "How do you define and measure the outcome? Who arbitrates disagreements?"

Figure 2. The Pricing-Positioning Connection. Pricing is not a finance decision with marketing inputs — it is a positioning, competitive, and psychological decision that requires PMM leadership.

The competitive intelligence dimension is particularly important. Your competitors are going through the same pricing transition you are, and their choices create opportunities and constraints. If your main competitor moves to consumption pricing and you stay on per-seat, you need to articulate why per-seat is better for the customer — and in some cases, it genuinely is.

Figure 2: Every pricing decision is three decisions at once.

The competitive intelligence dimension is particularly important. Your competitors are going through the same pricing transition you are, and their choices create opportunities and constraints. If your main competitor moves to consumption pricing and you stay on per-seat, you need to articulate why per-seat is better for the customer—and in some cases, it genuinely is. If your competitor offers AI credits and you don't, you need to either match the model or position against it. These are PMM problems.

They require the same kind of market intelligence, competitive analysis, and positioning skill that you bring to every other aspect of the go-to-market motion.

Agent-Powered Pricing Intelligence

The agentic tools from earlier chapters apply directly to the pricing function. An agent-powered competitive monitoring system can track pricing changes across your competitive landscape in near-real-time—new pricing pages, updated packaging tiers, customer reports of quoted prices, analyst commentary on pricing strategy. This intelligence is gold for the PMM who owns pricing inputs, because pricing decisions made without competitive context are pricing decisions made blind.

The agentic tools from earlier chapters apply directly to the pricing function. An agent-powered competitive monitoring system can track pricing changes across your competitive landscape in near-real-time — new pricing pages, updated packaging tiers, customer reports of quoted prices, analyst commentary on pricing strategy.

Figure 3: The Pricing Intelligence Pipeline

Three layers of agent-powered pricing intelligence. The signal layer runs continuously. The synthesis layer contextualizes. The response layer arms the team.

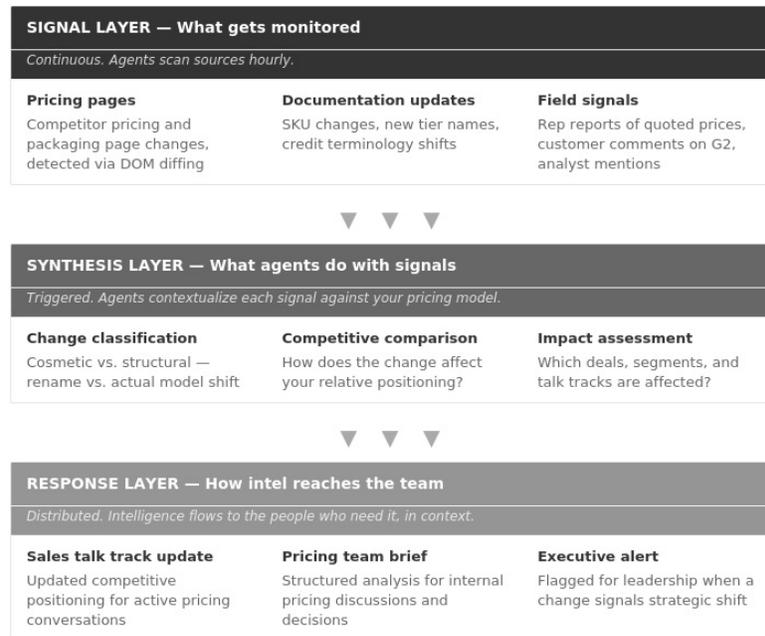


Figure 3. The Pricing Intelligence Pipeline. Without this pipeline, pricing decisions are made blind. With it, you caught the competitor's shift from "credit-based compute" to "AI workload units" within hours — not weeks.

Because we had the intelligence early, we had time to prepare a response — a clear articulation of how our consumption model compared, what our AI credit roadmap looked like, and how to frame the pricing conversation in competitive deals. Without the monitoring system, we would have been reactive. With it, we were proactive.

Figure 3: Three layers of agent-powered pricing intelligence.

We've seen this play out concretely. When a major competitor in the data platform space shifted their pricing from a capacity-based model to an AI-workload-based model, our monitoring system caught the change within hours. The initial signal was a documentation update—they revised their pricing page language from "credit-based compute" to "AI workload units." A few days later, an analyst mentioned in a brief that the competitor was repositioning their pricing narrative around "outcome-aligned costs." A week after that, we started hearing from sales reps that prospects were asking about our AI pricing model. Because we had the intelligence early, we had time to prepare a response—a clear articulation of how our consumption model compared, what our AI credit roadmap looked like, and how to frame the pricing conversation in competitive deals. Without the monitoring system, we would have been reactive. With it, we were proactive.

The Practitioner's Playbook: Pricing for the Agentic Era

Pricing is a chapter where the practitioner's playbook looks different depending on your role and your company's maturity, but there are some universal moves.

Pricing is a chapter where the practitioner's playbook looks different depending on your role and your company's maturity, but there are some universal moves.

Figure 4: The Pricing Practitioner's Playbook

Three moves for any PMM who wants to own the pricing conversation — regardless of whether you currently have a seat at the table.

01 Get in the Room

Lead with intelligence, not opinion

- The move** → Walk into the pricing meeting with data nobody else has: what competitors charge, how they've changed packaging, what buyers say about pricing expectations.
- Why it works** → Competitive pricing intelligence is your ticket in. Finance owns the model. Product owns the features. Nobody owns the market context — until you do.
- First step** → Pull the top three competitors' pricing pages. Build a comparison. Walk it into your next pricing discussion unsolicited.

02 Build the Pipeline

Continuous monitoring, not quarterly snapshots

- The move** → Add competitor pricing pages to your CI monitoring system. Track documentation changes as leading indicators of strategic shifts.
- Why it works** → Pricing shifts are signals. A competitor changing "credit-based compute" to "AI workload units" tells you something. An agent catches it in hours. You catch it in months.
- First step** → Add the pricing and packaging URLs for your top five competitors to your CI agent's watchlist. Set alerts for any DOM changes.

03 Own the Narrative

Translate mechanics into meaning

- The move** → Turn the pricing model into a story a rep can tell in a meeting without a spreadsheet. If they need a calculator, the narrative isn't ready.
- Why it works** → A consumption model that's well-explained feels transparent and fair. Poorly communicated, it feels unpredictable and risky. Same model, different outcome.
- First step** → Write a one-paragraph pricing story for your product: "You pay for X, here's how customers like you typically consume, here's what that costs." Test it on a rep.

The sales narrative test: If a rep can't explain your pricing model in a customer meeting without opening a spreadsheet, the model isn't ready for market — no matter how elegant the financial engineering.

Figure 4. The Pricing Practitioner's Playbook. Three moves that work regardless of company size, pricing maturity, or whether you currently have a seat at the pricing table. Start with Move 1. The rest follows.

How you talk about your pricing model to buyers is at least as important as the model itself. The PMM who can translate the financial mechanics of the pricing model into a buyer-friendly narrative — "you pay for what you use, and here's how customers like you typically consume the product" — is doing work that neither finance nor product can do as well.

Figure 4: Three moves for any PMM who wants to own the pricing conversation.

First, get in the room. If you're a PMM who isn't currently part of pricing decisions, make the case for inclusion by leading with competitive intelligence. The fastest way to get a seat at the pricing table is to walk in with data that nobody else has: here's what Competitor X is charging, here's how they've changed their packaging, here's what buyers are telling us about pricing expectations. That intelligence is your ticket in.

Second, build the pricing monitoring pipeline. Add competitor pricing pages to your CI monitoring system. Track documentation changes on pricing and packaging pages specifically—these are leading indicators of strategic shifts. Monitor review sites and forums for customer comments about pricing, which often surface ahead of official announcements. The goal is to ensure that your pricing team is never surprised by a competitive pricing move.

Third, own the pricing narrative. How you talk about your pricing model to buyers is at least as important as the model itself. A consumption model that's well-explained and well-positioned feels transparent and fair. The same model, poorly communicated, feels unpredictable and risky. The PMM who can translate the financial mechanics of the pricing model into a buyer-friendly narrative—"you pay for what you use, and here's how customers like you typically consume the product"—is doing work that neither finance nor product can do as well.

THE CMO PERSPECTIVE

Pricing is the area where the gap between what PMMs typically contribute and what the business needs is widest. The transition from per-seat SaaS to consumption-based pricing is the most significant business model shift in enterprise software since the on-prem-to-cloud migration. It changes the sales motion, the financial model, the competitive dynamic, and the PMM's job in ways that most PMMs aren't prepared for.

What the business needs from PMMs on pricing: continuous competitive pricing intelligence instead of quarterly snapshots. Customer research on willingness-to-pay and value perception—particularly around AI features, where pricing expectations aren't yet established. Packaging recommendations grounded in buyer segmentation and competitive positioning, not just financial optimization. And most importantly, a pricing narrative the sales team can actually deliver—because the most elegant pricing model is useless if the rep can't explain it without a spreadsheet. Pricing is positioning. The companies that get the AI pricing transition right will gain market share from those that don't.

KEY TAKEAWAYS

- **The per-seat-to-consumption shift** is the biggest enterprise software business model change since on-prem-to-cloud.
- **PMMs must own continuous competitive pricing intelligence**, not quarterly snapshots.
- **Every pricing model needs a "sales narrative test"**: if a rep can't explain it in a customer meeting without a spreadsheet, it's not ready.
- **Pricing is positioning**—get the AI pricing transition right and you gain share; get it wrong and you lose it.

CHAPTER NINE

Launch Management at Machine Speed

Pragmatic Remix: Go-to-Market Strategy • Launch Plan • Marketing Plan • Event Support

Chapter 9 Launch Management at Machine Speed Pragmatic Remix: Go-to-Market Strategy • Launch Plan • Marketing Plan • Event Support

The launch tracker was a Google Sheet with 247 line items. We know because we counted them at 11:30 PM on a Sunday, two days before a major product announcement, trying to figure out why the analyst briefing deck showed a different feature name than the press release. Somewhere between the product team's final naming decision on Thursday and the PR team's draft on Friday, the old name had persisted in one branch of the content tree while the new name propagated through the other. Nobody caught it because nobody could hold 247 line items in their head simultaneously, and the tracker—which was supposed to be the single source of truth—had been updated in the press release column but not the analyst materials column. If you've ever run a product launch at an enterprise software company, this story is painfully familiar. The specifics vary—maybe it's a pricing inconsistency rather than a naming one, or a screenshot that shows the old UI, or a customer quote that legal approved in version two but that got overwritten with an unapproved version in the final pass—but the underlying dynamic is always the same. Launch management is a coordination problem at scale, and the coordination complexity grows exponentially with the number of stakeholders, assets, channels, and time zones involved. A Tier 1 launch at a company like SAP might involve product marketing, product management, PR, analyst relations, demand generation, field marketing, sales enablement, partner marketing, legal, compliance, and executive communications—all producing and reviewing content on overlapping timelines, all working from a tracker that's perpetually three updates behind reality. This is Cluster One territory—operational coordination that is essential, timeconsuming, error-prone, and almost entirely automatable. And the automation isn't theoretical. It's the most straightforward agent application in the entire PMM job.

The Launch Coordination Agent The agent-powered launch workflow doesn't replace the launch. It replaces the tracker— and everything the tracker was supposed to do but couldn't. Imagine a system that ingests your launch plan—the list of deliverables, their owners, their dependencies, their review chains, and their deadlines—and actively manages the coordination. Not passively sitting in a spreadsheet waiting to be updated, but proactively monitoring: has the press release been reviewed by legal? Is the analyst deck consistent with the product page

copy? Has the sales enablement email been localized for EMEA? When the product team changes a feature name on Thursday, the agent identifies every downstream asset that references the old name, flags the inconsistency, and either proposes the updates or makes them automatically depending on the level of autonomy you've configured. This isn't a project management tool with AI features bolted on. It's a fundamentally different approach to launch coordination where the intelligence lives in the workflow rather than in the PMM's head. The PMM's job shifts from tracking the 247 line items to making the strategic decisions that the tracker can't make: Is the launch narrative strong enough? Are we timing this right relative to the competitive landscape? Is the tier classification correct, or should we upgrade this launch based on the market signal it sends? Those are judgment calls. The coordination is mechanical.

The Always-Launching Cadence There's a second transformation happening in launch management that's less about tools and more about cadence. The traditional launch model assumes that launches are events—discrete moments in time where the company announces something new, the marketing machine activates, and then everyone goes back to normal. Launches have tiers (Tier 1 gets the full treatment, Tier 3 gets a blog post and a prayer), and the cadence is typically quarterly, aligned with the product release cycle. The consumption-based pricing models from Chapter 8 are disrupting this cadence. When your business model is based on usage rather than contracts, the “launch” isn't the moment the customer signs; it's every moment the customer decides to use more.

Feature releases happen continuously. Capability improvements roll out weekly. The relevant marketing motion isn't a big-bang launch; it's a persistent drumbeat of value communication that keeps existing customers expanding and prospects engaged. This “always launching” cadence is impossible to execute manually at scale. A PMM who has to produce a launch brief, coordinate assets, brief sales, and update the website for every feature release would do nothing else. But an agent-powered launch pipeline can handle the Tier 3 and Tier 4 launches—the incremental updates, the feature improvements, the integration additions—almost autonomously. The agent generates the blog post from the product release notes. It updates the product page. It drafts the enablement email. It flags the update for the appropriate sales teams based on the feature's relevance to their active opportunities. The human PMM reviews the output, makes judgment calls about anything that needs strategic attention, and focuses their creative energy on the Tier 1 and Tier 2 launches where narrative, timing, and competitive positioning matter. The math is straightforward: if you can automate the 80% of launches that are operational and incremental, you free up massive capacity for the 20% that are strategic and career-making. That's the 10x move for launch management.

The Fast Bad Launch We want to spend a moment on what goes wrong, because the cautionary tale is as important as the optimistic case. The risk

of agent-powered launches isn't that they fail mechanically. The risk is that they succeed mechanically—that the assets get produced, the timelines get met, the coordination happens flawlessly—but the launch still fails because nobody applied judgment to the strategic questions. We've seen this happen when teams over-automate. The agent produces a launch brief that's technically correct but strategically wrong—it positions a defensive feature update as a Tier 1 innovation launch because the feature has broad applicability, not because it's strategically important. Or it generates a competitive narrative that's accurate but tone-deaf—aggressively calling out a competitor by name in a market where the buying committee values collaborative vendor relationships. The fast bad launch is the agentic era's version of the old software adage about shipping bugs faster. If your launch pipeline can go from product release to market announcement in 48 hours instead of four weeks, the cost of a bad strategic call gets amplified. You don't have four weeks of review cycles to catch the tone-deaf competitive positioning or the incorrect tier classification. The speed is a feature if the judgment is sound. It's a liability if it isn't. This is why the PMM's role in the agent-powered launch model is primarily editorial—in the best sense of the word. Not editing for grammar or formatting. Editing for strategy, tone, narrative coherence, and competitive awareness. The launch system proposes. The PMM disposes. And the PMM's ability to make fast, good judgment calls about what to amplify, what to tone down, what to delay, and what to kill becomes the most valuable skill in the launch motion.

The Practitioner's Playbook: Launch Management Three moves to make. First, audit your launch taxonomy. How many launches did your team execute last quarter? How many were Tier 1 versus Tier 3 or lower? If the answer is "we treat most launches the same because we don't have time to differentiate," you have a tier classification problem that agents can solve. Build a launch classification agent that evaluates each incoming product release against criteria you define—competitive significance, revenue impact, customer visibility, analyst relevance—and recommends a

tier. You review and adjust. The classification itself takes minutes instead of the two meetings and three Slack threads it currently requires. Second, automate the Tier 3 and Tier 4 pipeline. These are the incremental updates that need communication but don't need strategy. Build a workflow where the product release notes feed into an agent that generates the standard set of launch assets—blog post, product page update, enablement email, social post—and routes them for a quick human review before publication. This is the highest-leverage automation in the launch function because it reclaims the most time for the least risk. Third, invest the reclaimed time in launch narrative. For your Tier 1 and Tier 2 launches, spend the time you've freed up on the work that actually differentiates a great launch from a competent one: the opening story, the competitive framing, the analyst messaging, the executive talking points, the customer proof. These are the elements that turn a product announcement into a market moment, and they require the kind of

strategic creativity that no agent can provide.

Launch management is a coordination problem at scale, and the coordination complexity grows exponentially with the number of stakeholders, assets, channels, and time zones involved. A Tier 1 launch at a company like SAP might involve product marketing, product management, PR, analyst relations, demand generation, field marketing, sales enablement, partner marketing, legal, compliance, and executive communications.

If you can automate the 80% of launches that are operational and incremental, you free up massive capacity for the 20% that are strategic and career-making. That's the 10x move for launch management.

Figure 1: The Launch Tier Framework

Four tiers. The bottom two are automatable. The top two require judgment. Most teams treat them all the same — and that's the problem.

TIER	DESCRIPTION	ASSET SET	PMM ROLE	AUTOMATION
TIER 1 Bet-the-company	Major product launch, new category, strategic pivot. Shapes the narrative for the next 12 months.	Full launch: press, analysts, keynote, enablement, partner kit, exec comms, web, social, demand gen.	PMM leads strategy end-to-end. Narrative architect, competitive framing, timing.	Human-led Agents handle coordination only
TIER 2 Major	Significant feature, new integration, pricing change. Important but doesn't redefine positioning.	Blog, analyst brief, enablement email, product page update, sales deck addendum, social.	PMM owns narrative and competitive framing. Reviews all assets.	Hybrid Agents draft, humans edit for strategy
TIER 3 Minor	Feature improvement, capability extension, integration addition. Needs communication, not strategy.	Blog post, product page update, enablement email, social post.	PMM does quick review. 15-minute quality gate.	Agent-led Agent generates, human reviews
TIER 4 Maintenance	Bug fix, performance improvement, minor UX update. Needs documentation, not marketing.	Release notes, changelog entry, support article update.	No PMM involvement. Auto-published with notification.	Fully automated Agent publishes autonomously

Figure 1. The Launch Tier Framework. Tier classification determines resource allocation. Most teams default to treating everything as Tier 2 because they don't have time to differentiate — which means Tier 1 launches get under-invested and Tier 3/4 launches consume capacity they shouldn't.

The agent generates the blog post from the product release notes. It updates the product page. It drafts the enablement email. The human PMM reviews the output, makes judgment calls about anything that needs strategic attention, and focuses their creative energy on the Tier 1 and Tier 2 launches where narrative, timing, and competitive positioning matter.

Figure 1: The Launch Tier Framework

The agent-powered launch workflow doesn't replace the launch. It replaces the tracker — and everything the tracker was supposed to do but couldn't. Imagine a system that ingests your launch plan and actively manages the coordination. Not passively sitting in a spreadsheet waiting to be updated, but proactively monitoring.

Figure 2: Tracker vs. Launch Coordination Agent

The traditional launch tracker is passive. The launch coordination agent is active. The difference isn't tooling — it's where the intelligence lives.

Dimension	The Launch Tracker	The Launch Agent
Intelligence	Passive. A spreadsheet that records what humans tell it. Always three updates behind reality.	Active. Monitors assets, flags inconsistencies, proposes updates, routes approvals proactively.
Consistency check	Manual. Someone has to read every asset and compare. Nobody has time.	Automated. When a feature name changes Thursday, every downstream asset referencing the old name is flagged by Friday morning.
Status tracking	Self-reported. Owners update their own rows. "In progress" means anything from 10% to 90% done.	System-observed. Tracks actual document state — draft exists, legal has reviewed, localization complete.
Cross-functional coordination	Slack threads and email chains. "Did anyone update the analyst deck?" "I thought you were doing that."	Dependency-aware routing. Knows which assets depend on which inputs and alerts the right person at the right time.
Error detection	Late and human-dependent. Found at 11:30 PM on Sunday, two days before launch.	Early and systematic. Inconsistencies surfaced the moment they're introduced, not the moment a human notices.
PMM's time spent	60-70% on operational coordination. Tracking, chasing, reconciling, catching errors.	60-70% on strategic decisions. Narrative, timing, competitive framing, tier classification.

Figure 2. Tracker vs. Launch Coordination Agent. The PMM's job shifts from tracking the 247 line items to making the strategic decisions that no tracker — and no agent — can make.

The shift: The launch system proposes. The PMM disposes. The PMM's ability to make fast, good judgment calls about what to amplify, tone down, delay, or kill becomes the most valuable skill in the launch motion.

This isn't a project management tool with AI features bolted on. It's a fundamentally different approach to launch coordination where the intelligence lives in the workflow rather than in the PMM's head.

Figure 2: The Launch Coordination Agent workflow

The risk of agent-powered launches isn't that they fail mechanically. The risk is that they succeed mechanically — that the assets get produced, the timelines get met, the coordination happens flawlessly — but the launch still fails because nobody applied judgment to the strategic questions.

Figure 3: The Fast Bad Launch

Speed is a feature if the judgment is sound. It's a liability if it isn't. The fast bad launch is the agentic era's version of shipping bugs faster.

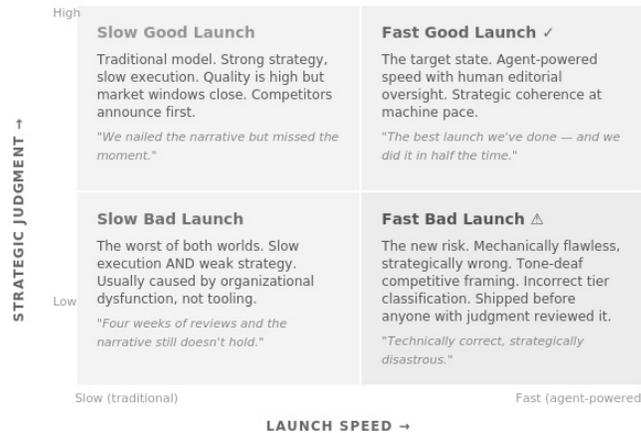


Figure 3. The Fast Bad Launch. Agent-powered pipelines move you from left to right on the speed axis. Only editorial judgment moves you from bottom to top. Without the editorial function, speed amplifies bad decisions.

This is why the PMM's role in the agent-powered launch model is primarily editorial — in the best sense of the word. Not editing for grammar or formatting. Editing for strategy, tone, narrative coherence, and competitive awareness.

Figure 3: The Launch Practitioner's Playbook

THE CMO PERSPECTIVE

Launch management is where the gap between expectation and execution is most frustrating—and it's not about talent. It's that the operational overhead of launches consumes so much energy that the strategic work gets squeezed. You can always tell when a launch was strategically underbaked: the announcement gets coverage but doesn't shift the narrative, enablement is technically correct but doesn't give reps a new story to tell, the analyst briefing checks boxes but doesn't move you in the evaluation. As teams gain more agent-powered tools, the corresponding investment should be in editorial review. We've found value in designating senior PMMs as "launch editors"—their job isn't to produce assets but to evaluate every launch for strategic coherence. Does the narrative match positioning? Is the competitive framing appropriate? Is the tier classification right? Is the timing smart relative to competitor activity? That editorial function is the quality gate that prevents speed from becoming recklessness. KEY TAKEAWAYS

- Operational overhead is the silent killer of launch quality—agent automation frees time for strategic coherence.
- The "fast bad launch" is a

real risk: speed without editorial review produces strategically underbaked announcements. • Designate senior PMMs as launch editors who evaluate narrative, competitive framing, and tier classification. • Launch tiers (bet-the-company, major, minor, maintenance) should dictate resource allocation, not politics.

Analyst Relations and the Influence Layer

Pragmatic Remix: Analyst Relations • Industry Relations • Influencer Relations

Chapter 10 Analyst Relations and the Influence Layer Pragmatic Remix:
Analyst Relations • Industry Relations • Influencer Relations

In 2024, one of our team leads showed us a spreadsheet tracking the time we spent on analyst RFIs—the detailed questionnaires that firms like Gartner, Forrester, and IDC send when they’re evaluating vendors for Magic Quadrants, Waves, MarketScapes, and the rest of the industry’s competitive evaluation apparatus. The numbers were staggering. A single Gartner Magic Quadrant RFI consumed, on average, 120 personhours across product marketing, product management, and engineering. We were completing six to eight of these per year across different product lines. That’s somewhere between 720 and 960 hours annually—the equivalent of half a full-time employee doing nothing but answering analyst questions. The RFIs are important. Gartner alone influences more than \$100 billion in enterprise software purchasing decisions annually, and a Magic Quadrant placement can make or break a sales quarter. But the process of completing them is brutally inefficient. Each RFI asks dozens of detailed questions about product capabilities, customer references, market strategy, competitive positioning, and roadmap direction. Many of the questions are similar across different evaluations—the Gartner BI Magic Quadrant and the Gartner Planning Magic Quadrant ask overlapping questions about data integration, AI capabilities, and cloud architecture—but the answers need to be customized for each evaluation’s specific criteria and context. We built an agent to fix this. Not to replace the human judgment that goes into crafting an RFI response—the strategic decisions about what to emphasize, how to position roadmap items, which customer references to cite—but to handle the mechanical retrieval and first-draft generation that consumed most of those 120 hours. The agent draws from a curated knowledge base of past RFI responses, product documentation, customer success stories, and competitive positioning documents. It generates a firstpass response for each question, pulling from the most relevant source material and flagging where the previous response may be stale or where new product capabilities need to be incorporated. The result: we reduced RFI response effort by 50 to 60 percent, and in some cases by as much as 70 percent on targeted workflows where the questions closely matched previous evaluations. The quality didn’t suffer—in fact, it improved, because the agent

was more consistent about pulling in the latest product capabilities and customer proof points than a human working against a deadline. The human’s job became editing, strategic adjustment, and the genuinely hard

work of deciding how to position emerging capabilities that hadn't been part of previous evaluations. • • •

The Analyst Landscape Is Shifting The RFI automation story is important, but it's the easy part. The harder and more interesting transformation in analyst relations is about what the analysts themselves are becoming—and how the influence landscape is fragmenting in ways that change the PMM's job. The big three—Gartner, Forrester, IDC—still dominate enterprise technology purchasing decisions. That's not changing anytime soon. But the way they work is changing. Analysts are processing more vendor information, faster, using their own AI tools. A Gartner analyst evaluating twenty vendors for a Magic Quadrant is almost certainly using AI to help synthesize the mountains of RFI data, briefing notes, and customer reference feedback they receive. This means that the structured, specific, evidencebacked communication that works for buyer agents (Chapter 4) also works for analyst evaluations. Vague claims get filtered out. Specific capabilities with quantified outcomes get weighted. At the same time, new influence channels are emerging. Independent analysts and practitioners with large followings—people building audiences on Substack, LinkedIn, and YouTube—are shaping buyer perceptions in ways that didn't exist five years ago. When a respected practitioner writes a detailed comparison of data platforms based on their actual experience, that content gets shared in Slack channels and buying committee discussions alongside the Gartner report. It's a different kind of authority—experiential rather than institutional—but it's increasingly influential. And then there's the AI influence layer—the GEO dimension from Chapter 6. When a buyer asks Perplexity or ChatGPT to recommend data platform vendors, the AI's response is shaped by the same content that shapes analyst opinions: product documentation, review sites, published evaluations, customer case studies, and vendor content. This isn't analyst relations in the traditional sense, but it's influence—and managing it requires many of the same skills: understanding what the evaluating entity is looking for, ensuring your story is told accurately and compellingly in the channels that matter, and maintaining the consistency and evidence depth that builds credibility.

The PMM as Influence Architect The traditional AR function is relatively narrow: manage the relationship with Gartner, Forrester, and IDC analysts who cover your category. Prepare for briefings. Respond to RFIs. Lobby for positioning. Provide customer references. Track placements. In the agentic era, the AR function expands into something broader that we'd call influence architecture: the deliberate design and management of how your product is perceived across all the evaluation channels that matter to your buyers. This includes traditional analysts, but it also includes independent voices, AI systems, peer communities, and the increasingly important layer of customer advocates who share their experiences publicly. The skill set shifts accordingly. The traditional AR skill was relationship management—knowing the analyst, understanding their evaluation criteria, building rapport over time. That skill still matters. But the influence architect also

needs to think about structured information delivery (how to make your story evaluable by AI systems and AI-augmented analysts), content strategy (how to ensure the right content exists in the right channels to be surfaced by independent voices and AI search), and evidence management (maintaining a continuously updated library of customer proof points, quantified outcomes, and capability evidence that can be deployed across any influence channel on demand). The agent stack from earlier chapters supports this directly. Your competitive monitoring system tracks not just competitor moves but analyst commentary and sentiment shifts. Your knowledge base feeds both RFI responses and GEO-optimized content. Your content pipeline produces material that serves human analysts, AI systems, and buyer communities simultaneously. The influence architect orchestrates all of these into a coherent strategy that ensures your product's story is told accurately and compellingly wherever buyers look for evaluation input.

The Practitioner's Playbook: Analyst Relations and Influence If you own AR or influence at your company, here's where to focus. First, build the RFI knowledge base. Even if you don't automate the full RFI workflow, creating a curated repository of past responses, organized by topic and tagged with recency and accuracy flags, will cut your response time significantly. Every time you

complete an RFI, add the final responses to the knowledge base with metadata about which evaluation it was for, which analyst, and when. Over two to three evaluation cycles, you'll have a corpus that an agent can draw from effectively. Second, map your influence landscape. Make a list of every entity that shapes how buyers in your category evaluate vendors: the major analyst firms, the independent analysts and practitioners, the review sites (G2, TrustRadius, Gartner Peer Insights), the AI systems that surface vendor recommendations, the community forums and Slack groups where practitioners share opinions. For each entity, assess your current presence: are you well-represented? Is the information accurate and current? Are there gaps? This map becomes your influence strategy—a prioritized list of channels where investment will move the needle on buyer perception. Third, make the analyst briefing a two-way intelligence channel. Most PMMs treat analyst briefings as a pitch—an opportunity to tell the analyst your story. The better play is to use briefings as intelligence gathering. What are analysts hearing from buyers? What evaluation criteria are shifting? Which competitive narratives are gaining traction? What are they skeptical about in your story? An analyst who trusts you enough to share candid feedback is giving you intelligence that's worth more than the Magic Quadrant placement itself, because you can use it to improve your positioning, your product, and your competitive strategy.

A single Gartner Magic Quadrant RFI consumed, on average, 120 person-hours across product marketing, product management, and engineering. We were completing six to eight of these per year. That's the equivalent of half a full-time employee doing nothing but answering analyst questions.

We built an agent to fix this. Not to replace the human judgment that goes into crafting an RFI response — but to handle the mechanical retrieval and first-draft generation that consumed most of those 120 hours.

Figure 1: The RFI Automation Impact

Agent-powered RFI workflow: automate retrieval and first-draft generation, keep strategic judgment human.

120 Person-hours per Magic Quadrant RFI	6-8 RFIs completed per year	720-960 Hours annually ≈ 0.5 FTE on RFIs	50-70% Reduction in response effort
---	---------------------------------------	---	---

Dimension	Agent Handles	Human Owns
Source retrieval	Draws from curated knowledge base of past responses, product docs, customer stories.	Curates and maintains the knowledge base. Adds new responses after each cycle.
First-draft generation	Generates first-pass response for each question from relevant sources.	Reviews, edits, strategically adjusts. Decides emphasis and positioning.
Staleness detection	Flags where previous responses are stale or new capabilities need incorporation.	Decides how to position emerging capabilities not in previous evaluations.
Consistency	More consistent at pulling latest capabilities and proof points under deadline.	Strategic judgment on which customer references to cite for this evaluation.

Quality didn't suffer — it improved. The agent was more consistent about pulling in the latest capabilities than a human working against a deadline. The human's job became editing, strategic adjustment, and positioning emerging capabilities.

Figure 1. The RFI Automation Impact. 120 person-hours per RFI, reduced by 50-70%. The human's job shifts from retrieval and drafting to editing and strategic positioning.

The result: we reduced RFI response effort by 50 to 60 percent, and in some cases by as much as 70 percent on targeted workflows where the questions closely matched previous evaluations.

Figure 1: The AR Influence Model

The big three — Gartner, Forrester, IDC — still dominate enterprise technology purchasing decisions. But the influence landscape is fragmenting in ways that change the PMM's job.

Figure 2: The Influence Landscape

Four layers of influence that shape how buyers evaluate vendors. Traditional AR covers one. The influence architect covers all four.

LAYER 1: Traditional Analysts		
<i>Institutional authority. Structured evaluations. Still dominant.</i>		
Who Gartner, Forrester, IDC — the big three who evaluate vendors for Magic Quadrants, Waves, MarketScapes.	What's changing Analysts are using AI to synthesize vendor data. Vague claims get filtered. Specific, quantified capabilities get weighted.	PMM implication Structured, evidence-backed communication matters more than ever. The same discipline that works for buyer agents (Ch 4) works here.
LAYER 2: Independent Voices		
<i>Experiential authority. Practitioner credibility. Increasingly influential.</i>		
Who Practitioners with large followings on Substack, LinkedIn, YouTube. They write from actual experience, not vendor briefings.	What's changing Their content gets shared in Slack channels and buying committee discussions alongside the Gartner report. Different authority — experiential, not institutional.	PMM implication You can't brief them like analysts. You earn their attention by building a product worth writing about and relationships worth maintaining.
LAYER 3: AI Influence Layer		
<i>Algorithmic authority. GEO-driven. The fastest-growing evaluation channel.</i>		
Who Perplexity, ChatGPT, Claude, Gemini — AI systems buyers query for vendor recommendations.	What's changing AI responses are shaped by the same content that shapes analyst opinions: product docs, review sites, evaluations, case studies.	PMM implication GEO optimization (Ch 6) becomes an AR discipline. Your content must be structured for AI consumption and discoverable across evaluation channels.
LAYER 4: Peer Communities		
<i>Social proof. Review sites. Community consensus.</i>		
Who G2, TrustRadius, Gartner Peer Insights, community Slack groups, Reddit, practitioner forums.	What's changing Peer reviews are now a weighted signal for both human buyers and AI systems. Low review scores get surfaced in AI recommendations.	PMM implication Customer advocacy programs feed all four layers. A strong review presence improves analyst perception, AI visibility, and peer credibility simultaneously.
The multiplier: A Leader placement in the Magic Quadrant isn't just a website badge — it's a data point that AI systems weight heavily when filtering vendor shortlists. Every influence layer feeds the others.		

Figure 2. The Influence Landscape. Four layers of evaluation influence. Traditional AR covers Layer 1. The influence architect covers all four — because buyers, analysts, and AI agents all draw from the same evidence ecosystem.

This isn't analyst relations in the traditional sense, but it's influence — and managing it requires many of the same skills: understanding what the evaluating entity is looking for, ensuring your story is told accurately in the channels that matter.

Figure 2: The Agent-Powered AR System

The traditional AR function is relatively narrow: manage the relationship with Gartner, Forrester, and IDC analysts. Prepare for briefings. Respond to RFIs. Lobby for positioning. Track placements. In the agentic era, the function expands into something broader: influence architecture.

Figure 3: The Influence Architect

The skill set shift from traditional AR manager to influence architect. The relationship skill still matters — but it's now one of four capabilities.

Capability	Traditional AR	Influence Architect
Relationship management	Know the analyst. Build rapport. Understand their evaluation criteria. Manage the briefing calendar.	Still essential — but now extends to independent voices, community moderators, and understanding how AI systems weight different sources.
Structured information delivery	Write clear RFI responses. Prepare polished briefing decks.	Make your story evaluable by AI systems and AI-augmented analysts. Structured, specific, evidence-backed across every channel.
Content strategy	Produce analyst-facing content: briefing slides, RFI responses, reference stories.	Ensure the right content exists in the right channels to be surfaced by independent voices, AI search, and peer communities simultaneously.
Evidence management	Maintain a list of customer references for analyst requests.	Continuously updated library of customer proof points, quantified outcomes, and capability evidence deployable across any influence channel on demand.

The orchestration model: Your competitive monitoring system tracks analyst commentary and sentiment shifts. Your knowledge base feeds RFI responses AND GEO-optimized content. Your content pipeline serves human analysts, AI systems, and buyer communities simultaneously. The influence architect orchestrates all of these into coherent strategy.

Figure 3. The Influence Architect. From AR manager to influence architect. The relationship skill is now one of four capabilities required to manage how your product is perceived across all evaluation channels that matter to buyers.

The influence architect orchestrates competitive monitoring, knowledge management, content strategy, and evidence management into a coherent strategy that ensures your product's story is told accurately and compellingly wherever buyers look for evaluation input.

Figure 3: AR Maturity Framework

If you own AR or influence at your company, here's where to focus. Each move compounds — the knowledge base feeds the influence map, and both feed the briefing strategy.

Figure 4: The AR Practitioner's Playbook

Three moves that transform analyst relations from a Gartner relationship into a full-spectrum influence strategy.

01

Build the RFI Knowledge Base

The foundation everything else draws from

- The move** → Create a curated repository of past RFI responses, organized by topic, tagged with recency and accuracy flags. Every time you complete an RFI, add the final responses with metadata: which evaluation, which analyst, when.
- Why it works** → Over two to three evaluation cycles, you'll have a corpus that an agent can draw from effectively. Even without full automation, this cuts response time significantly and eliminates the "start from scratch every time" problem.
- First step** → Collect your last three completed RFIs. Tag each response by topic (product capabilities, AI, cloud architecture, security, customer references). That's your seed corpus.

02

Map Your Influence Landscape

Every entity that shapes buyer evaluation

- The move** → List every entity that shapes how buyers evaluate vendors: major analyst firms, independent analysts, review sites (G2, TrustRadius, Peer Insights), AI systems, community forums and Slack groups. For each, assess: are you well-represented? Is the info current? Are there gaps?
- Why it works** → This map becomes your influence strategy — a prioritized list of channels where investment will move the needle on buyer perception. Most AR teams manage Layer 1 and ignore the other three.
- First step** → Ask Perplexity to recommend vendors in your category. See what comes back. That's your AI influence audit — and it's often a wake-up call.

03

Make Briefings a Two-Way Intelligence Channel

The briefing as intelligence gathering, not just pitching

- The move** → Stop treating analyst briefings as one-way pitches. Use them to gather intelligence: what are analysts hearing from buyers? What criteria are shifting? Which competitive narratives are gaining traction? What are they skeptical about in your story?
- Why it works** → An analyst who trusts you enough to share candid feedback is giving you intelligence worth more than the placement itself — because you can use it to improve positioning, product, and competitive strategy.
- First step** → In your next analyst briefing, allocate the last 15 minutes to questions: "What are you hearing from buyers about [our category]? What do you think we're missing?"

The reframe: AR ROI in the agentic era includes a multiplier — AI agents weight analyst evaluations when filtering vendor shortlists. A Gartner Leader placement isn't just credibility for human buyers. It's a high-weighted signal for every AI agent that helps a buyer make a shortlist. Think influence holistically.

Figure 4. The AR Practitioner's Playbook. Three moves: build the knowledge base that powers RFI automation, map the full influence landscape beyond the big three, and transform analyst briefings from one-way pitches to two-way intelligence channels.

The traditional model treats analyst relations as a tax — something you have to do to maintain placements. In the agentic era, that framing is even less appropriate. The ROI of AR isn't just the human buyers who read the report. It's the AI agents that cite the report.

Figure 4: The AR Practitioner's Playbook

Analyst relations is one of the areas where agent-powered workflows have delivered the most tangible ROI—and the RFI automation described in this chapter is a real example from our team. But the broader point is about reframing AR investment entirely. The traditional model treats analyst relations as a tax—something you have to do to maintain Gartner and Forrester placements, a cost center that doesn't directly generate pipeline. In the agentic era, that framing is even less appropriate. When a buyer's agent evaluates vendors, it draws on analyst evaluations as a high-credibility signal. A Leader placement in the Magic Quadrant isn't just a website badge—it's a data point that AI systems weight heavily when filtering vendor shortlists. The ROI of AR isn't just the human buyers who read the report. It's the AI agents that cite the report. Think about influence holistically—analyst placements plus GEO, review sites, independent voices, and the full spectrum of evaluation channels. The PMM who thinks about influence this way will outperform the PMM who treats AR as a Gartner relationship. KEY TAKEAWAYS

- AR ROI in the agentic era includes a multiplier: AI agents weight analyst evaluations when filtering vendor shortlists.
- Reframe AR as a discoverability investment, not just a credibility investment.
- RFI automation dramatically reduces the operational burden that makes AR feel like a tax.
- Think influence holistically: analyst placements, GEO, review sites, and independent voices all feed agent evaluation.

PART III Building the 10x PMM Organization

The PMM Tech Stack: A Practitioner's Evaluation

Pragmatic Remix: Sales Tools • Marketing Plan • Competitive Landscape • Content Creation

Chapter 11 The PMM Tech Stack: A Practitioner's Evaluation Pragmatic Remix: Sales Tools • Marketing Plan • Competitive Landscape • Content Creation

In late 2025, we put together an AI tools acquisition plan for our PMM organization at SAP—fifty-three people across product marketing, competitive intelligence, pricing, and research. The exercise was supposed to take a week. It took three, because the landscape is a mess. We don't mean the tools are bad. Many of them are excellent. We mean that the number of AI tools claiming to solve PMM problems has exploded past the point where any individual can evaluate them rationally. At last count, there were over 200 tools in the "AI for marketing" category on G2 alone—and that doesn't include the general-purpose LLMs, the developer-focused agent platforms, or the enterprise AI suites that include marketing features as part of a larger offering. A PMM trying to build a coherent tech stack is facing the same problem that CMOs faced with martech a decade ago: too many tools, too many categories, too many vendor claims, and not enough clarity about what actually works in practice. This chapter is our attempt to cut through that noise. Not with a comprehensive market map—those go stale before the ink dries—but with a practitioner's framework for thinking about which tools matter, which categories are real, and how to build a stack that's coherent rather than a Frankenstein of point solutions that each solve one problem while creating three new integration headaches.

The Core Stack Every PMM needs a foundation layer, and that foundation is a general-purpose LLM assistant. Full stop. If you don't have a Claude, ChatGPT, or Gemini subscription—or if your company hasn't provisioned one for you—fix that before you read another word of this chapter. The general-purpose LLM is the Swiss Army knife of the agentic PMM. It's your first-draft generator, your research synthesizer, your brainstorming partner, your copy editor, your data analyst, and your format converter. It's not the best tool for any

single one of those tasks, but it's good enough at all of them that it covers 70% of what a PMM needs to get started. When we did the evaluation for our team, we recommended Claude as the primary LLM for a few reasons that we think generalize beyond our specific context. Claude's strength in long-form writing, nuanced analysis, and maintaining context across complex multistep tasks makes it particularly well-suited to PMM work, where the typical output isn't a quick answer but a positioning

document, a competitive analysis, or a strategic recommendation that requires holding multiple dimensions in mind simultaneously. ChatGPT has strengths in breadth and multimodal capabilities—the image generation and web browsing are genuinely useful—and Gemini has advantages in Google ecosystem integration. But for the core PMM workflow of research, analysis, and writing, Claude was the strongest fit. Alongside the LLM, we recommended two other core tools. Grammarly Pro—not for grammar checking, which the LLM handles fine, but for brand voice consistency and tone management across a fifty-person team producing content in multiple channels. And Perplexity Enterprise for research—because Perplexity’s web-grounded, citationbacked research workflow is better than any general-purpose LLM for the kind of competitive intelligence and market research that PMMs need to do quickly and accurately. That’s the core: Claude for creation and analysis, Grammarly for consistency, Perplexity for research. Three tools. Everything else is additive.

The Specialist Layer On top of the core stack, there are category-specific tools that are worth evaluating based on your team’s specific needs. we’ll cover the categories we’ve evaluated most thoroughly. Competitive intelligence platforms—Klue, Crayon, Kompyte—are purpose-built for the continuous monitoring and battlecard workflows we described in Chapters 3 and 5. They offer out-of-the-box competitor tracking, automated alerting, and integration with sales enablement platforms. The advantage over building your own CI pipeline from scratch is speed to value and maintenance—these platforms handle the plumbing so you can focus on the intelligence. The disadvantage is cost (enterprise pricing starts around \$30K annually) and flexibility—you’re limited to the platform’s data sources and synthesis capabilities, which may not match what you can build with a custom agent pipeline. For teams at earlier maturity levels or with smaller budgets, the custom CI pipeline approach we outlined in Chapter 3—RSS monitoring plus LLM synthesis plus a structured delivery mechanism—can achieve 80% of what a dedicated platform provides at a fraction of the cost. The trade-off is maintenance: you’re building and supporting the infrastructure yourself. Content generation platforms—Jasper, Writer, Copy.ai—are specialized for marketing content production. They offer brand voice templates, campaign workflows, and team collaboration features that general-purpose LLMs don’t provide natively. We tested these against Claude for our content workflows and found that for commodity content (the bottom layer from Chapter 6), they offer a workflow advantage: the templating and brand consistency features save time when you’re producing high volumes of channelspecific content. For strategic and signature content, the general-purpose LLM is better because it handles the nuance, context, and analytical depth that content platforms sacrifice for workflow efficiency. Demo automation platforms—the full landscape from Chapter 7—are worth evaluating separately from the AI stack because they solve a different class of problem. But they’re

increasingly incorporating AI features (auto-personalization,

engagement analytics, conversational demo experiences) that make them part of the agent-powered PMM workflow. Research tools beyond Perplexity—Elicit for academic and technical research, Consensus for evidence synthesis—are useful for PMMs who do deep market analysis, pricing research, or product strategy work that requires grounding in quantitative evidence.

The Build-vs-Buy Decision The most important decision in building your PMM tech stack isn't which tool to buy. It's whether to buy a specialized tool or build a custom workflow using the generalpurpose LLM and agent platforms. We've seen this play out across our team and across the PMM teams we advise. The instinct—especially for PMMs who don't have a technical background—is to buy. Specialized tools have UIs, onboarding flows, customer success teams, and the psychological comfort of a vendor relationship. The instinct for technically inclined PMMs is to build. Agent platforms like LangChain and CrewAI offer enormous flexibility, and there's a genuine thrill in building a custom system that does exactly what you need. The right answer is usually a combination, and the deciding factor is maintenance burden. A CI monitoring pipeline that you build yourself is great—until the RSS feeds break and nobody fixes them, or the LLM's output format changes and the parsing logic fails, or you leave the company and nobody understands how the system works. Purpose-built platforms handle that maintenance. The trade-off is flexibility and cost. My general guidance: buy for workflows that are stable and well-defined (competitive monitoring, content management, demo automation), where the platform's maintained infrastructure is worth the cost. Build for workflows that are experimental or unique to your organization (custom CI synthesis, specialized RAG systems, pricing intelligence pipelines), where the flexibility of a custom approach is worth the maintenance investment. And start with the general-purpose LLM for everything, because it's the

fastest way to learn which workflows are worth investing in—you can always specialize later.

Security, Governance, and the Enterprise Reality We need to address something that every PMM at a large enterprise will encounter: the security and governance conversation. When you propose adopting AI tools for your team, your IT and security organizations are going to have questions—legitimate ones— about data handling, access controls, model training, and compliance. The most common concern we hear is about data leakage: if PMMs are putting competitive intelligence, pricing strategies, product roadmaps, and customer data into AI tools, where does that data go? Does the AI vendor use it to train their models? Could a competitor's employee, using the same tool, access intelligence that your team provided? The answer depends entirely on which tier of service you're using. Consumer-tier LLM subscriptions typically include data in model training unless you opt out. Enterprise tiers—Claude Team, ChatGPT Enterprise, Gemini for Workspace—typically include contractual guarantees that customer data isn't used for training and that access controls meet enterprise standards. The cost difference between consumer

and enterprise tiers is significant, but for a PMM team handling competitive intelligence and pricing strategy, the enterprise tier isn't optional. It's the cost of doing business responsibly. Our advice: don't wait for IT to come to you. Go to them with a proposal that addresses their concerns proactively. Specify which tools you want, at which tier, with which data handling guarantees. Show that you've done the homework on security and compliance. In our experience, the PMMs who get AI tool adoption approved fastest are the ones who frame it as a responsible business case, not a request for permission to experiment.

The Practitioner's Playbook: Building Your Stack Here's the recommendation by team maturity level. If you're a solo PMM or a team of two to three, start with the core stack: one enterprisetier LLM (Claude or ChatGPT), Perplexity for research, and whatever content management tool your company already uses. Invest your time in learning the LLM

deeply—building custom prompts for your recurring workflows, experimenting with context window management for complex analyses, and developing a personal RAG system for your competitive and product knowledge. The general-purpose LLM, used well, can cover 90% of what a small team needs. If you're a team of five to fifteen, add the specialist layer: a competitive intelligence platform if your competitive landscape is complex enough to warrant it, a content generation tool if your production volume is high, and a demo automation platform if your product supports self-serve evaluation. At this scale, the integration between tools starts to matter—make sure your CI platform feeds your enablement workflow, your content tool draws from your messaging framework, and your analytics span across tools. If you're a team of fifteen-plus, you're building an ecosystem, and the biggest risk is fragmentation. Designate a PMM operations role—or at least a PMM who owns the stack—to ensure coherence. Build the connective tissue between tools: shared knowledge bases, consistent taxonomies, unified analytics. And invest in custom agent pipelines for the workflows that are unique to your organization, because at this scale the general-purpose tools won't cover your specific needs.

The number of AI tools claiming to solve PMM problems has exploded past the point where any individual can evaluate them rationally. Over 200 tools in the "AI for marketing" category on G2 alone. A PMM trying to build a coherent stack is facing the same problem CMOs faced with martech a decade ago.

That's the core: Claude for creation and analysis, Grammarly for consistency, Perplexity for research. Three tools. Everything else is additive.

Figure 1: The PMM Tech Stack Architecture

Three layers. Start at the bottom. Add upward only when you've exhausted the layer below. The core covers 70% of what a PMM needs.



Figure 1. The PMM Tech Stack Architecture. Three layers, bottom-up. The core stack (3 tools) covers 70% of PMM needs. The specialist layer adds category-specific capability. The custom layer provides maximum flexibility at maximum maintenance cost.

This chapter is our attempt to cut through that noise. Not with a comprehensive market map — those go stale before the ink dries — but with a practitioner's framework for building a stack that's coherent.

Figure 1: The PMM Tech Stack

The number of AI tools claiming to solve PMM problems has exploded past the point where any individual can evaluate them rationally. Over 200 tools in the "AI for marketing" category on G2 alone. A PMM trying to build a coherent stack is facing the same problem CMOs faced with martech a decade ago.

That's the core: Claude for creation and analysis, Grammarly for consistency, Perplexity for research. Three tools. Everything else is additive.

Figure 1: The PMM Tech Stack Architecture

Three layers. Start at the bottom. Add upward only when you've exhausted the layer below. The core covers 70% of what a PMM needs.



Figure 1. The PMM Tech Stack Architecture. Three layers, bottom-up. The core stack (3 tools) covers 70% of PMM needs. The specialist layer adds category-specific capability. The custom layer provides maximum flexibility at maximum maintenance cost.

This chapter is our attempt to cut through that noise. Not with a comprehensive market map — those go stale before the ink dries — but with a practitioner's framework for building a stack that's coherent.

Figure 2: Tech Stack Integration Points

The most important decision in building your PMM tech stack isn't which tool to buy. It's whether to buy a specialized tool or build a custom workflow using general-purpose LLM and agent platforms. The deciding factor is maintenance burden.

Figure 2: The Build vs. Buy Decision

When to buy a platform, when to build custom, and when to start with the general-purpose LLM. The deciding factor isn't capability — it's maintenance burden.

Workflow	Recommendation	Rationale	Trade-off
Competitive monitoring	BUY if complex landscape START WITH LLM if simple	Stable, well-defined workflow. Platform handles RSS feeds, parsing, alerting. Maintained infrastructure worth the cost.	Cost (~\$30K/yr) vs. flexibility. Platform's data sources may not match what you'd build custom.
Content production	BUY for commodity volume START WITH LLM for strategic	Content platforms offer brand voice templates and campaign workflows the LLM doesn't have natively. Good for high-volume commodity content.	Quality vs. workflow efficiency. Strategic/signature content is better with the general-purpose LLM.
Demo automation	BUY	Solves a different class of problem. Increasingly AI-native. Can't replicate with an LLM alone.	Build economics. Weeks to set up, ongoing maintenance as product ships new features.
CI synthesis & analysis	BUILD	Experimental, unique to your org. Custom agent pipeline gives you exactly the synthesis you need.	Maintenance is on you. RSS feeds break. LLM output formats change. Parsing logic fails.
RFI knowledge base	BUILD	Your past responses + your product docs + your competitive positioning = unique corpus. No vendor has this.	Requires curation discipline. Stale entries degrade quality over time.
Pricing intelligence	BUILD	Custom pipeline from Ch 8. Your competitors, your pricing model, your specific dimensions. No platform covers this well.	Technically ambitious. Needs someone who can maintain the pipeline long-term.
Everything else	START WITH LLM	The general-purpose LLM is the fastest way to learn which workflows are worth investing in. Specialize later.	You can always add layers — it's harder to untangle a Frankenstein of point solutions.

The deciding factor: Maintenance burden. Buy for workflows that are stable and well-defined, where maintained infrastructure is worth the cost. Build for workflows that are experimental or unique, where flexibility is worth the maintenance investment. Start with the LLM for everything else.

Figure 2. The Build vs. Buy Decision. The right answer is usually a combination. Buy for stable workflows (CI monitoring, content, demo). Build for unique workflows (custom synthesis, knowledge bases, pricing intel). Start with the LLM for everything else.

Start with the general-purpose LLM for everything, because it's the fastest way to learn which workflows are worth investing in — you can always specialize later.

Figure 3: Build vs Buy Decision Matrix

The stack recommendation scales with team size — not because bigger teams need fancier tools, but because the integration and coherence challenges multiply as the team grows.

Figure 3: The Stack by Team Size

What to deploy at each scale. The core stack is the same for everyone. What changes is the specialist layer and the need for connective tissue between tools.

1-3 <small>people</small>	Solo / Small Team <i>The LLM covers 90% of what you need</i>	
Tools Core stack only; enterprise-tier LLM + Perplexity + existing content management.	Investment focus Learn the LLM deeply. Build custom prompts for recurring workflows. Develop a personal RAG for competitive and product knowledge.	Key principle Depth over breadth. One tool used expertly beats five tools used casually.
5-15 <small>people</small>	Mid-Size Team <i>Add the specialist layer selectively</i>	
Tools Core + CI platform (if complex landscape) + content generation (if high volume) + demo automation (self-serve eval).	Investment focus Integration between tools. CI feeds enablement. Content draws from messaging framework. Analytics span across tools.	Key principle Integration over accumulation. Three connected tools beat six siloed ones.
15+ <small>people</small>	Large Team / Ecosystem <i>The biggest risk is fragmentation</i>	
Tools Core + specialist + custom agent pipelines for workflows unique to the org. Full three-layer stack.	Investment focus Connective tissue: shared knowledge bases, consistent taxonomies, unified analytics. Designate a PMM ops role to own the stack.	Key principle Coherence over capability. A team with 12 AI tools and no coherent workflow is paying tool tax, not getting leverage.
The fragmentation risk: The worst outcome is a team with twelve AI tools and no coherent workflow, where the PMM spends half their time copying information between systems. That's not leverage — that's tax. Simplify before you add.		

Figure 3. The Stack by Team Size. The core is universal. The specialist layer scales with need. The custom layer adds flexibility at maintenance cost. At every scale, coherence beats capability.

Designate a PMM operations role — or at least a PMM who owns the stack — to ensure coherence. Build the connective tissue between tools.

Figure 4: The Tech Stack Practitioner's Playbook

THE CMO PERSPECTIVE

The tools conversation is the one PMM teams bring to leadership most frequently, and here's how it should be evaluated—because understanding the CMO's framework helps any PMM make a better case. Three criteria matter. First, leverage:

does this tool enable one person to produce what previously required two, or the same output fifty percent faster? Second, quality: does it maintain or improve output quality, or trade quality for speed? A tool that makes you faster at lower quality is a bad investment because the downstream cost of bad intelligence or positioning is higher than the time savings. Third, coherence: does it fit the existing workflow or create another silo? The worst outcome—and it happens—is a team with twelve AI tools and no coherent workflow, where the PMM spends half their time copying

information between systems. That's not leverage; that's tax. When making the case for a tool, the strongest signal isn't the vendor's pitch deck. It's a demonstration that you've already tried it, built a workflow around it, and can show the output versus what you were producing before. Don't ask for permission to evaluate. Evaluate first and come with evidence. KEY TAKEAWAYS

- Evaluate tools on three criteria: leverage (output per person), quality (maintained or improved), and coherence (workflow integration).
- A small core stack plus selective specialist tools beats a dozen disconnected point solutions.
- The strongest tool pitch isn't a vendor deck—it's a before/after comparison from your own workflow.
- Avoid the "tool tax": if your team spends more time moving data between tools than using it, the stack needs simplifying.

Skills That Survive: What AI Can't Touch (Yet)

Pragmatic Remix: Distinctive Competence • Innovation Games • Market Problems • Buying Process

Chapter 12 Skills That Survive: What AI Can't Touch (Yet) Pragmatic Remix: Distinctive Competence • Innovation Games • Market Problems • Buying Process

The question we get asked most often—at conferences, in DMs, over drinks at industry dinners—is some version of: “Should I be worried about my job?” The PMMs who ask it range from people two years into their career to directors with fifteen years of experience. The junior ones are worried about whether there will be a PMM role to grow into. The senior ones are worried about whether the role they’ve spent a decade mastering is about to be pulled out from under them. Both groups deserve an honest answer, so here it is: no, your job is not going away. But the version of your job you’re doing today probably is. That’s not a hedge. It’s a precise distinction. The activities that define “product marketing”—the thirty-seven boxes on the Pragmatic Framework—aren’t disappearing. Companies still need competitive intelligence, positioning, sales enablement, content, launch management, pricing strategy, and analyst relations. They need them more than ever, actually, because the agentic era is making go-to-market more complex, not less. What’s disappearing is the specific way those activities have been performed: manually, sequentially, at human speed, with human bottlenecks at every step. The PMM of 2023 spent the majority of their time on execution—producing deliverables. The PMM of 2028 will spend the majority of their time on judgment—making decisions that agents can’t make. The deliverables still get produced. They just get produced differently, at different speed, and the human’s contribution shifts from drafting to directing. So the real question isn’t “Will I have a job?” It’s “Which skills do I need to develop so that the version of me that shows up in 2028 is more valuable, not less?” That’s what this chapter is about.

The Four Surviving Skills When we did the Pragmatic Framework analysis in Chapter 2, the activities that scored highest on elevation opportunity—the ones that become more valuable as the routine work around them is automated—clustered around four capabilities. These aren’t job titles or Pragmatic boxes. They’re meta-skills that cut across multiple activities and that define what it means to be a senior, strategic product marketer in the agentic era.

Strategic Judgment Strategic judgment is the ability to make good decisions under ambiguity with incomplete information. It sounds abstract until you realize it’s the core of almost every Cluster Three activity. Which

competitive move do you respond to and which do you ignore? Your CI agent flags twelve competitive signals this week. Three of them are noise. Six are worth noting but don't require action. Two require immediate response. One requires a fundamental rethinking of your competitive narrative. The agent can't tell you which is which. It can prioritize by recency, by signal strength, by relevance to your stated competitive priorities—but the judgment call about whether a competitor's new partnership is a strategic threat or a desperate pivot requires understanding the competitive landscape, the competitor's internal dynamics, your own company's strategic position, and what the market will perceive. That's judgment. It's developed through experience, not training. Which pricing model do you adopt when the data is inconclusive? Your pricing research shows that 40% of customers prefer consumption-based pricing and 35% prefer per-seat. The remaining 25% say it depends on the use case. There is no clear winner. An agent can present the data. It can model scenarios. It can show you what competitors are doing. But the decision to go with consumption pricing for your enterprise tier and per-seat for your SMB tier, with a conversion path between them—that's a strategic judgment that synthesizes market data, competitive positioning, sales team capabilities, and financial model constraints into a single coherent choice.

Strategic judgment is hard to develop deliberately because it's mostly developed through exposure to decisions and their consequences. But there are some accelerants: postmortems on every major decision (not just the ones that went wrong), exposure to decision-making in other functions (sit in on product roadmap reviews, pricing committee meetings, sales pipeline reviews), and the habit of writing down your reasoning before you see the outcome so you can evaluate your judgment calibration over time.

Organizational Intelligence We mentioned in Chapter 2 that getting twelve people to agree on a positioning document is harder than writing it. That's organizational intelligence—the ability to navigate stakeholders, build cross-functional alignment, manage up and across, and understand the informal power structures that determine what actually gets done. This is the skill that AI is furthest from replicating, and we think it's the one that separates good PMMs from great ones. A good PMM can write excellent positioning. A great PMM can get that positioning adopted by product, endorsed by sales, supported by leadership, and reflected in every customer-facing touchpoint. The writing is maybe 20% of the effort. The organizational navigation—the meetings, the compromises, the one-on-one conversations where you understand what someone really objects to (which is usually not what they say they object to), the political awareness of when to push and when to wait—is the other 80%.

We'll give you a specific example because abstract advice about "navigating stakeholders" is useless without context. Early in our time at SAP, we were working on positioning for a product that sat at the intersection of two business units. Each business unit had a VP who wanted the positioning to emphasize their product's contribution. The

positioning document went through seven drafts—not because the writing was

wrong but because each draft surfaced a new political objection that was really about organizational territory, not about messaging. We solved it by having coffee with each VP separately, not to discuss the positioning but to understand what they needed the positioning to accomplish for their team. One VP needed the positioning to support a headcount request. The other needed it to justify a product investment to the board. Once we understood what each person actually needed—as opposed to what they said they wanted in the positioning document—we could write a version that served both agendas without compromising the market-facing clarity. No agent is going to have that coffee. No agent is going to read the body language that tells you the real objection isn't about the word choice on slide seven. This skill is developed by paying attention to people—their incentives, their anxieties, their ambitions—and by building enough trust that they'll tell you what they actually think rather than what they think you want to hear.

Narrative Craft We covered this in Chapter 6, so we'll keep it brief here. Narrative craft isn't the same as writing skill. Agents can write well. Narrative craft is the ability to construct an argument that moves people—that takes a complex, abstract, technical reality and turns it into a story that someone remembers the next day.

The skill has three components. First, story selection: knowing which story to tell for which audience at which moment. The Keurig chip story works in a chapter about DMP economics because it makes data collection feel personal and specific. It would be wrong in a chapter about pricing because it's not about pricing. The ability to match stories to

arguments is a form of taste that's developed through wide reading, careful observation, and a lot of practice. Second, specificity: the discipline of using concrete details rather than abstract claims. Named people. Named companies. Specific numbers. The exact question Douwe Bergsma asked in the conference room. The twelve-year-old sneaking Donut Shop KCups on weekends. These details are what make a story sticky, and they require having been somewhere and paid attention. Third, voice: the quality that makes your writing recognizably yours. Not performative quirks or stylistic tics, but the authentic expression of how you see the world and what you think is funny, important, and true. This is the hardest component to develop because it requires enough self-knowledge to write honestly and enough confidence to write distinctively. Most PMMs default to a corporate voice that's competent and invisible. The ones who break through have figured out how to sound like themselves.

Customer Empathy We put customer empathy last not because it's least important but because it's the one most PMMs think they already have—and most of them are wrong, or at least incomplete, about what it actually requires.

Customer empathy isn't the same as customer knowledge. An agent can give you comprehensive customer knowledge: usage patterns, support tickets, NPS scores, feature requests, churn predictors. That's data. Customer empathy is the ability to understand what a customer feels—the frustration of trying to get budget approval for a tool that everyone agrees they need but nobody will prioritize, the anxiety of betting

their reputation on a vendor choice that could go wrong, the relief when a product actually solves the problem they were hired to solve. The PMM who has customer empathy produces different work than the PMM who has customer knowledge. The knowledge-based PMM writes positioning that maps features to requirements. The empathy-based PMM writes positioning that maps capabilities to emotions: “Stop dreading Monday’s forecast meeting” hits differently than “Improve forecast accuracy by 34%”—not because the second claim is wrong but because the first one connects with the human experience of what bad forecasting actually feels like in a planning team’s daily life. Customer empathy is developed through direct customer exposure: interviews, ridealongs, user testing sessions, support call listening, and—critically—informal conversations where you’re not asking structured research questions but just talking to a person about their job. We make it a rule to have at least two unstructured customer conversations per month, and we encourage every PMM we manage to do the same. The insights from these conversations don't always show up in a deliverable, but they accumulate into a felt understanding of the customer that informs everything you produce.

The T-Shaped PMM The model that brings all of this together is what we think of as the T-shaped PMM. The horizontal bar of the T is broad competence across agent-augmented skills: the ability to direct agents for competitive intelligence, content production, data analysis, launch coordination, and the other Cluster One and Cluster Two activities. Every PMM needs this bar. It's the table stakes of the agentic era—the minimum level of agent fluency required to remain competitive.

The vertical bar of the T is deep expertise in one of the four surviving skills: strategic judgment, organizational intelligence, narrative craft, or customer empathy. This is your differentiator—the skill that makes you irreplaceable rather than merely competent. The PMM with deep narrative craft and broad agent fluency is a different archetype than the

PMM with deep organizational intelligence and broad agent fluency, and both are valuable in different contexts. The career implication is clear: invest in both bars. Develop your agent fluency broadly—learn the tools, build the workflows, get comfortable directing rather than doing. And simultaneously deepen your strongest surviving skill to the point where you're recognized for it—the person people come to for the hard strategic calls, or the person who can navigate any stakeholder minefield, or the person whose writing stops people in their feed, or the person who understands customers at a level that everyone else is just guessing at. The PMM with both bars—broad agent fluency and deep human skill—is

the 10x PMM. Not because they work harder than their peers, but because they've figured out what to outsource to machines and what to keep for themselves. That distinction is, in a sense, the entire argument of this book.

replicate on its own. These are not the places where AI helps you go faster. They are the places where AI-generated output without expert human direction produces something generic. The Full Stack PMM directs the horizontal and owns the vertical.

The AI-Enabled Workflow

The T-shape describes what the Full Stack PMM knows. The workflow describes how they actually spend their time. The gap between the PMM still doing the work the old way and the one who has redesigned around AI is most visible not in the tools they use but in how the day is allocated.

The following comparison is constructed from conversations with PMMs who have made the transition — not the ones who bought tools and declared victory, but the ones who redesigned their workflows and are now operating differently. They did not get more hours. They did not hire more people. They reallocated the same eight hours.

Figure 2: The AI-Enabled Workflow

BEFORE: Traditional PMM Day		→	AFTER: Full Stack PMM Day	
Morning brief prep Scanning news, compiling competitive updates	20%		AI-supervised execution Agents produce drafts; PMM reviews and directs	10%
First draft content Battlecards, one-pagers, blog posts, briefs	30%		Competitive intelligence Always-on agent monitoring; PMM interprets and acts	22%
Research & synthesis Win/loss reviews, analyst reports, competitor sites	20%		Narrative & positioning More time here than ever — the irreplaceable work	28%
Internal meetings Syncs, reviews, stakeholder updates	15%		Strategic & exec work Roadmap influence, exec comms, category bets	25%
Strategic work Positioning, roadmap input, exec narrative	15%		Stack management Directing agents, improving prompts, QA-ing output	15%

Figure 2. Same eight hours, fundamentally different allocation. Commodity execution drops from ~50% to ~10%. Narrative and strategic time more than doubles.

The reallocation in the after column is not aspirational. It is what the data shows in organizations that have genuinely integrated AI into their workflow. The commodity execution category collapses from roughly fifty percent to roughly ten. That freed capacity goes to narrative and positioning, strategic and executive work, and competitive intelligence

Figure 1: The Four Surviving Skills

that is now continuous and agent-monitored rather than quarterly and manual.

The most counterintuitive item in the after column is the growth in narrative and positioning time. Most PMMs assume that AI will help them produce more content faster — and it does. But the Full Stack PMMs who have made the transition report spending more time on positioning and narrative than before, not less. The commodity production work was consuming time they should have been spending on thinking through the story. AI gave them that time back.

The Depth/Breadth Matrix

The T-shape describes the Full Stack PMM profile. The Depth/Breadth Matrix describes the landscape of profiles that exist in the real world — and the dynamics that determine which ones compound over time and which ones fall behind.

Figure 3: The Depth/Breadth Matrix

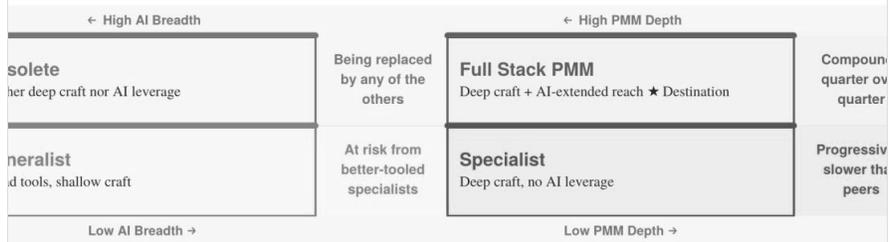


Figure 3. Four profiles. The destination is top-right: deep craft + AI reach. The risk is the top-left: obsolete before you know it.

The four quadrants map to profiles that show up in every PMM organization. Most PMMs sit in the bottom two: Specialist (deep but slow) or Generalist (fast but shallow). Both are viable in the short term. Neither compounds. The Specialist's depth is real but increasingly slow to deploy at scale. The Generalist's speed is real but increasingly indistinguishable from what any decently-prompted AI can produce.

The Obsolete quadrant is the one most PMM leaders are not yet honest about. It is populated not by people who are behind on AI adoption but by people who added AI tools to an unstructured workflow and did not develop genuine domain expertise. The tools did not save them. The lack of depth made the tools produce generic output. For some portion of the current PMM population, the Obsolete quadrant is already the present.

Figure 2: The AI-Enabled Workflow

The Stack Architecture

The Full Stack PMM operates across three layers, each with a distinct function and a distinct relationship to the PMM's own judgment. Understanding the layers is useful not just as a conceptual model but as an operational guide — it clarifies which work the PMM owns, which work they direct, and which work can be delegated to agents entirely.

Figure 4: The Stack Architecture

3 OUTPUT LAYER <i>What the market sees</i> Human judgment owns this layer entirely. These outputs define the PMM's value.	Positioning docs	Competitive battlecards	Exec narratives
	Launch assets	Category POVs	
2 SPECIALIST LAYER <i>Domain intelligence and synthesis</i> Agents produce the raw material. The PMM directs, filters, and interprets.	Win/loss synthesis	Competitor monitoring	Voice-of-customer
	Market trend tracking		
1 CORE AI LAYER <i>Automation and velocity</i> Fully AI-executed. PMM sets standards and reviews output.	First-draft generation	Format conversion	Research aggregation
	Prompt library		

Figure 4. Three layers: Core AI (velocity), Specialist (intelligence), Output (judgment). The PMM owns all three but executes only on the top. The judgment loop runs upward from each layer.

The judgment loop is the element of the stack architecture that gets the least attention but produces the most leverage. It is the mechanism by which the PMM's insight at the Output Layer improves everything produced at the layers below. When a PMM reviews output and decides it is not sharp enough — the positioning is generic, the competitive analysis missed the angle — that judgment feeds back into the prompts and frameworks that govern the Specialist Layer agents. Over time, the stack gets better because the PMM keeps improving the judgment loop. This is the compounding mechanism that the Full Stack PMM has and the AI Generalist does not.

Three Moves

Figure 3: The Skills Development Path

The second failure mode is the AI Generalist: broad tool fluency, shallow craft. This PMM has built an impressive personal stack. They know how to prompt, chain tools, automate. But their actual domain expertise is thin — their positioning is generic, their competitive analysis is surface-level, their narrative lacks the depth that earns executive trust. AI amplifies what you bring. If what you bring is generic, AI makes you generically faster.

The two traps are mirrors of each other. Pure Specialists have depth but not leverage. AI Generalists have leverage but not depth. The Full Stack PMM escapes both.

Figure 1: The Full Stack PMM T-Shape

AI-ENABLED HORIZONTAL REACH				
Research & Intelligence	Content Production	Demo & Enablement	Data & Analytics	Comms & Distribution
Rapid synthesis, always-on CI	First draft auto, format scaling	AE readiness, asset updates	Attribution, conversion intel	Message testing, channel optim.
		Positioning		
		Competitive Intelligence		
		Narrative		
		Category Design		
<i>AI handles the horizontal. Same output quality at a fraction of the time. The Full Stack PMM directs, not executes.</i>			<i>The depth is irreplaceable. AI replicates breadth; it cannot replicate judgment. The vertical is where value compounds.</i>	

Figure 1. The T-shape: depth in core PMM craft (vertical) + AI-extended reach across five adjacent functions (horizontal). The depth is irreplaceable; the horizontal is AI-enabled.

The T-shape is not a new concept — it has been used in product and design to describe deep expertise extended by broad awareness. What makes the Full Stack PMM T-shape specific to the AI era is the mechanism. The horizontal breadth is not achieved by generalist experience. It is achieved by deploying AI agents across five adjacent functions and directing them with enough domain judgment to produce usable output.

The vertical depth is where the Full Stack PMM's irreplaceable value lives. Positioning, competitive intelligence, narrative construction, and category design are the activities where deep PMM craft produces outcomes that AI cannot

Figure 4: The Skills Practitioner's Playbook

THE CMO PERSPECTIVE

When hiring for 10x PMMs, the most revealing interview question may be: "Tell me about a time you changed your mind about something important based on new information. What was the original position, what changed it, and what did you do differently?" The answer reveals nearly everything that matters. Whether they can be specific rather than abstract—because the best answers are stories with named people, real stakes, and concrete outcomes. Whether they have strategic judgment— because changing your mind professionally requires intellectual honesty and clarity about what "right" looks like. Whether they have organizational intelligence— because changing direction usually means convincing others too. What this question doesn't ask about is tools, frameworks, or deliverables. Those things matter, but they're learnable. The four skills that survive— judgment, organizational intelligence, narrative craft, and empathy—are the hard ones. They're what separates PMMs who will thrive in the agentic

era from those who will struggle. And they're what CMOs should be hiring for.

KEY TAKEAWAYS

- The four durable PMM skills: strategic judgment, organizational intelligence, narrative craft, and customer empathy. • These are the skills that agents can't replicate—and the ones that will define career trajectories in the agentic era. • Interview for process and adaptability, not just deliverables: "how do you work?" reveals more than "what have you shipped?" • Technical fluency with AI tools is table stakes; the differentiator is critical thinking about what the tools produce.

Team Structure and Hiring for the Agentic Org

Pragmatic Remix: All 37 Activities • Organizational Design • Talent Strategy

Chapter 13 Team Structure and Hiring for the Agentic Org Pragmatic Remix: All 37 Activities • Organizational Design • Talent Strategy

This chapter is different from the others, and we want to be transparent about why. For most of this book, we've been writing as practitioners—PMM leaders talking to PMMs about how to do the job better. This chapter leans more heavily on the executive view, because the questions it addresses—how to structure a team, who to hire, how to manage performance—are fundamentally organizational questions. They require the perspective of someone building and running a fifty-plus-person PMM team in real time, and that perspective shapes everything that follows.

The Org Chart Question Here's a thought experiment that captures the organizational tension perfectly. In late 2025, we were looking at our org chart—fifty-three people across product marketing, competitive intelligence, pricing and commercialization, and research—and running a thought experiment that every CMO is quietly running: if I were building this team from scratch today, with full access to agentic tools, would it look like this?

The honest answer was no. Not because the people were wrong—the team was strong— but because the structure was optimized for a pre-agentic operating model. The team was organized by function: a competitive intelligence group, a content and messaging group, a technical marketing group, a pricing and commercialization group. Each function produced its own deliverables, maintained its own tools, and operated on its

own cadence. The structure made sense when the bottleneck was execution—when producing a competitive battlecard required a different skill set from producing a pricing analysis, and each required enough specialized effort to justify a dedicated team. In the agentic model, the bottleneck shifts from execution to strategy. The competitive battlecard and the pricing analysis are both produced by agents drawing from a shared knowledge base. The human work—the judgment about what the competitive landscape means for pricing, or what the pricing shift means for competitive positioning—cuts across functional boundaries. The functional structure doesn't just feel inefficient; it actively prevents the cross-pollination of insight that makes the strategic work valuable.

From Role-Based to Outcome-Based The organizational shift we're implementing at SAP—and that we think is directionally correct for most

PMM teams—is from role-based teams to outcome-based teams. In a role-based structure, teams are organized around what they produce: the competitive team produces competitive content, the content team produces marketing content, the enablement team produces sales tools. Each team has a defined scope, clear deliverables, and specialized expertise. The advantage is clarity. The disadvantage is that the most important PMM work—the strategic synthesis that connects competitive intelligence to positioning to pricing to sales enablement—falls between teams rather than within them. In an outcome-based structure, teams are organized around what they achieve: a team owns pipeline generation for a specific product line, or competitive win rate for a specific segment, or analyst perception for the portfolio. Each team has the full spectrum of PMM capabilities—competitive, content, enablement, pricing—and the agent stack to execute across all of them. The team doesn't produce battlecards and positioning documents as separate deliverables; it produces a coherent go-to-market narrative that manifests as battlecards, positioning, sales enablement, and content simultaneously, because the same intelligence and the same strategic framework inform all of them.

This isn't just an org design preference. It's a response to the reality that agentic tools make cross-functional execution possible for small teams in a way that it wasn't before. A team of three PMMs with a coherent agent stack can cover competitive intelligence, positioning, enablement, and content for a product line—work that previously required six to eight specialized roles—because the agents handle the execution while the humans handle the strategy and judgment. The team is smaller not because the work went away but because the leverage per person increased.

The Full-Stack PMM The individual contributor who thrives in the outcome-based model is what we've been calling the full-stack PMM—someone who combines strategic depth, agent orchestration capability, and execution judgment across the full scope of PMM activities.

This is a different profile from the specialist PMM who has historically been the ideal hire. The specialist—the competitive intelligence expert, the pricing analyst, the content strategist—was valuable because their deep expertise in one area produced higher-quality deliverables than a generalist could. In the agentic era, the agent handles much of the specialized execution, and the human's value comes from connecting insights

across domains. The PMM who can see the thread from a competitive pricing shift to a positioning opportunity to a sales enablement need to an analyst briefing talking point—and who can direct agents to execute across all four—is more valuable than four specialists each seeing their piece of the picture. We don't want to overclaim here. There are still roles where deep specialization matters—particularly in pricing and commercialization, where the analytical complexity requires genuine expertise, and in analyst relations, where relationship depth takes years to build. But for the majority of PMM roles—the people who own a product

line's go-to-market—the full-stack profile is where the market is heading.

Hiring for the Agentic Era If you're hiring PMMs—or if you're a PMM thinking about how to position yourself in the market—here's what the evaluation criteria look like. Systems thinking matters more than deliverable experience. The question isn't "have you produced a competitive battlecard?" It's "have you designed a competitive

intelligence workflow?" The shift from artifact to system is the defining move of the agentic era, and PMMs who think in systems—inputs, processes, outputs, feedback loops—will navigate it better than PMMs who think in deliverables. Intellectual curiosity matters more than tool proficiency. The tools are going to change. Claude today won't be Claude in two years. The platforms from Chapter 11 will consolidate, evolve, or be replaced. What endures is the disposition to experiment, to learn new tools quickly, to build workflows iteratively, and to be honest about what's working and what isn't. The PMM who has tried six tools and has an informed opinion about each is a better hire than the PMM who has mastered one tool and never looked at the others. Communication range matters more than communication polish. The agentic era PMM needs to communicate with product engineers about data architecture, with sales reps about deal strategy, with analysts about market evaluation, with executives about business performance, and with agents about task execution. Each audience requires a different register. The PMM who can code-switch across these audiences—who can be technical with engineers and narrative with executives without sounding fake in either context—has a communication range that's more valuable than flawless copywriting. Adaptability matters more than pedigree. The best predictor of success in the agentic era isn't whether someone has a Pragmatic certification or ten years of PMM experience. It's whether they've demonstrated the ability to adapt when the rules changed. Did they learn new skills when their company shifted from on-prem to cloud? Did they evolve their approach when their market shifted from enterprise to mid-market? Did they experiment with AI tools before being told to? The pattern of adaptation is what you're hiring for, because the rules are about to change again.

The Performance Measurement Problem One of the hardest management challenges in the agentic era is measuring PMM performance when the traditional output metrics—number of battlecards produced, content pieces published, launches executed—become meaningless. If an agent produces the battlecard, what did the PMM contribute?

We've talked about this extensively, and we think the answer is to shift from output metrics to impact metrics. Output metrics measure what was produced. Impact metrics measure what changed as a result. Did the competitive win rate improve? Did the sales cycle shorten? Did the analyst move us up in the evaluation? Did the pricing change increase deal size? Did the launch shift the market narrative? Impact metrics are harder to measure and harder to attribute. They require longer time horizons and more tolerance for ambiguity. But they're the metrics that actually matter

—the ones that tell you whether the PMM is creating strategic value rather than just producing artifacts. And in the agentic era, where artifact production is cheap, strategic value is the only defensible basis for headcount. The practical approach: define two to three impact metrics for each PMM role at the beginning of the quarter. Competitive PMM: win rate in competitive deals and timefrom-signal-to-sales-response. Launch PMM: pipeline generated from launch activities and analyst sentiment shift. Pricing PMM: deal size growth and pricing-related win/loss trends. These are imperfect metrics—influenced by factors beyond the PMM’s control— but they point the team toward the right kind of work, which is the point of measurement.

were, in the vast majority of cases, describing a state that amounted to license procurement and individual experimentation. They were not describing structural change.

The author heard some version of the same conversation with every PMM leader consulted for this chapter. 'We have ChatGPT. Some of our team uses Claude. We got a Perplexity license. We're covered.' None of them could describe what 'covered' meant in operational terms. None had changed how they staffed commodity work. None had a definition of readiness that extended beyond the procurement decision. That is adoption. It is not integration.

Tool procurement is not a strategy. It is a precondition for a strategy — one that most organizations have mistaken for the strategy itself.

Figure 1: The PMM Readiness Spectrum

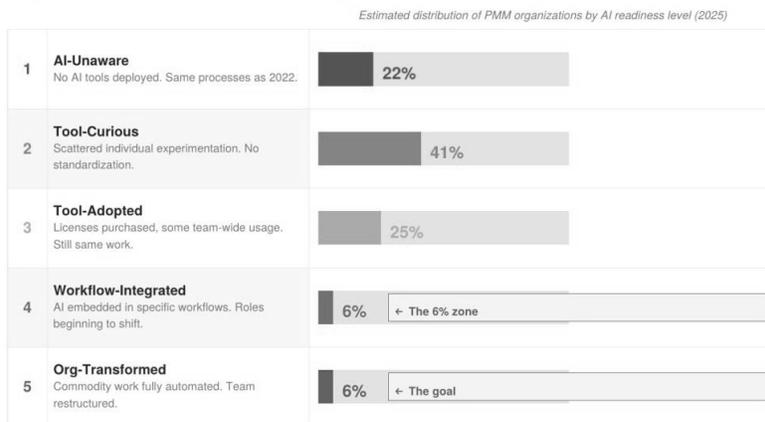


Figure 1. Five levels of AI integration, from AI-Unaware to Org-Transformed, with estimated distribution across marketing organizations. The 63% at Levels 1–2 have either no tools or only scattered individual usage. The 25% at Level 3 have purchased tools but not changed how the organization works. Only the 6% at Levels 4–5 are genuinely integrating.

The distribution is not surprising once you understand the structural incentive. Adopting AI tools is easy — it requires a purchase order and an announcement. Integrating AI requires answering questions that are much harder: which roles need to change, how do we measure the impact, what does it mean to succeed, and — most importantly — what are we going to stop doing so that we have capacity to do the irreplaceable work that actually compounds? Those questions do not have easy answers, and most PMM leaders are not yet asking

Figure 1: The Agentic PMM Team Structure

them explicitly enough.

Adoption vs. Integration

The adoption-integration distinction is not semantic. It describes a real and measurable difference in organizational behavior — one that shows up in how work is structured, how success is measured, and how the PMM team spends its time. The following framework makes the distinction precise enough to be diagnostic.

Figure 2: Adoption vs. Integration

DIMENSION	ADOPTION	INTEGRATION
How AI is used	<i>Faster way to do the same work</i>	Structural change to what work is done and by whom
Primary benefit	<i>Speed on existing tasks</i>	Capacity shift from commodity to irreplaceable work
Workflow change	<i>Individual, ad hoc, optional</i>	Team-wide, standardized, expected
Staffing implication	<i>Same headcount, same roles</i>	Roles redesigned; capacity redeployed
Content architecture	<i>Same formats, faster to produce</i>	Agent-readable formats added; GEO signals built in
Success metric	<i>'We're using AI more'</i>	'Commodity hours down X%; strategic hours up X%'
Risk	<i>Complacency disguised as adoption</i>	Organizational change requires leadership to lead it

Most organizations are at adoption. The 6% are at integration. The gap is decisions, not tools.

Figure 2. Seven dimensions on which adoption and integration differ. The integration column describes what is structurally different in organizations at Levels 4–5 of the readiness spectrum. The adoption column describes what most organizations believe counts as progress.

The most diagnostic dimension is the success metric. Organizations at the adoption level measure AI progress by usage — license activation, prompt volume, tools deployed. Organizations at the integration level measure it by the shift in where their people's time goes: what percentage of PMM hours are now on commodity versus irreplaceable work, how has that ratio changed quarter over quarter, and what is the quality impact on the strategic outputs that matter? The measurement distinction is not just operational — it is a leading indicator of whether the organization is building compounding leverage or just making the same work slightly faster.

Figure 2: Role Evolution Matrix

The staffing implication row deserves specific attention, because it is where most organizations are most evasive. If AI integration has not changed how you staff commodity work — if your PMM organization has the same roles doing the same mix of activities it had before the AI era — then the integration is not real. It is adoption with better PR. The PMMs who were spending sixty percent of their time on commodity work are still spending sixty percent of their time on commodity work, just faster. That is not leverage. That is efficiency applied to the wrong problem.

The Six Percent Diagnostic

The following eight questions are designed to expose the gap between what an organization believes about its AI readiness and what is operationally true. They are not a survey — they are a confrontation. Every PMM leader who answered them honestly in the research for this chapter arrived at a lower score than they expected.

Figure 3: The Six Percent Diagnostic

DIAGNOSTIC QUESTION	ADOPTION ANSWER	INTEGRATION ANSWER
1. Has your PMM org explicitly mapped which activities are now AI-executable?	<i>No formal mapping</i>	Yes — and capacity has been redeployed
2. Do you have standardized prompts/agents for your highest-volume PMM tasks?	<i>Individual experimentation only</i>	Team-wide library, maintained and updated
3. Has AI changed how you staff commodity work (briefs, research, first drafts)?	<i>Same headcount, same roles</i>	Roles redesigned; fewer people needed for commodity
4. Do you measure AI productivity impact — hours recovered, quality delta?	<i>No measurement framework</i>	Tracked and reported to leadership
5. Has your content architecture changed to include agent-readable formats?	<i>Same content types, faster production</i>	Schema markup, structured specs, GEO signals added
6. Do your competitive assets update continuously rather than quarterly?	<i>Quarterly manual update cycle</i>	Agent-monitored, near-real-time updates
7. Can your PMM team articulate the difference between adoption and integration?	<i>Not part of team vocabulary</i>	Active distinction, used in planning
8. Has leadership set integration targets with named accountability?	<i>AI as individual initiative</i>	Org-level KPIs; leader-owned

Figure 3. Eight diagnostic questions that distinguish AI-adopted from AI-integrated organizations. The adoption answer describes what most organizations can truthfully claim. The integration answer describes what the 6% have actually built.

The diagnostic is most useful when applied as a team exercise rather than a self-assessment. The gap between what a PMM leader believes is true about their organization and what individual contributors believe is true is

Figure 3: Team Capacity Model

often illuminating — and the illumination is almost always in the same direction. Leaders overestimate the degree of integration. ICs underestimate the urgency. Scoring the diagnostic together, with honest answers from both levels of the org, tends to produce clarity that no strategy document achieves.

Every PMM leader who answered these questions honestly in the research for this chapter arrived at a lower score than they expected. That gap between perception and reality is precisely the problem this chapter is trying to name.

The Org Design Shift

The hardest part of the six percent problem is not the diagnostic — it is what the diagnostic implies. If AI integration is real, it does not produce a PMM organization that is the same size doing more work. It produces a PMM organization that is structured differently, staffed differently, and optimized for a fundamentally different mix of activities. The org design shift is not optional; it is the logical endpoint of genuine integration.

Figure 4: The Org Design Shift

BEFORE: Traditional PMM Org		AFTER: AI-Integrated PMM Org	
60–70 %	Commodity execution Content, research, templates, first drafts	~5%	AI-supervised commodity Agents produce; PMMs review and direct
15–20 %	Enablement ops Battlecard maintenance, training, RFPs	20–25 %	Competitive intelligence Living systems, agent-monitored, always-on
10–15 %	Campaigns support Brief writing, copy, asset management	25–30 %	Narrative & positioning More time here than ever before
~10%	Strategic work Positioning, narrative, category, exec comms	40–50 %	Strategic & exec work Positioning bets, roadmap influence, exec comms

Figure 4. Before and after comparison of PMM org time allocation. The traditional org spends 60–70% of capacity on commodity execution. The AI-integrated org inverts that ratio — with commodity work reduced to ~5% under AI supervision, and 40–50% of capacity reallocated to strategic and executive work.

The inversion in the 'after' column is not aspirational. It is what integration actually looks like when the commodity work — content drafts, competitive research, template production, basic enablement — flows through AI agents rather than through human PMMs. The PMMs who were doing that work do not disappear; they become the directors and editors of AI output, spending their judgment on the things that require judgment

Figure 4: The Team Structure Playbook

THE CMO PERSPECTIVE

The organizational questions this chapter raises are the ones that should keep every CMO productively restless—because the team structure decisions made in the next twelve months will determine organizational effectiveness for the next five years. What we’re seeing work: moving from a purely functional structure to a hybrid that maintains centers of excellence for deep specialization (pricing, analyst relations) while reorganizing the rest around product-line outcomes. Each product line gets a small team of two to four PMMs who own the full go-to-market, supported by shared CoEs and the agent infrastructure the team builds. On headcount: we are not using agentic tools to cut staff. We’re using them to reallocate time from low-value execution to high-value strategy. The PMMs who used to spend sixty percent of their time producing artifacts now spend sixty percent on competitive analysis, customer insight, positioning refinement, and stakeholder navigation. Output quality

went up. Strategic impact went up. And the ability to justify investment went up, because PMM activity now connects to pipeline and revenue metrics that leadership cares about. On hiring: we now include a “work with AI” exercise, we’ve deprioritized portfolio reviews in favor of strategic decision stories, and we’ve made customer interaction experience a hard requirement. The framing that captures it: we’re hiring architects, not builders. Agents handle the building. We need people who understand the full system. KEY TAKEAWAYS

- Hybrid team structures—centers of excellence for deep specialization plus product-line outcome teams—are emerging as the model.
- Agentic tools should reallocate time from execution to strategy, not reduce headcount.
- New hiring signals: “work with AI” exercises, strategic decision narratives over portfolio reviews, and mandatory customer experience.
- The organizing metaphor: hire architects who understand the full system, not builders. Agents handle the building.

CONCLUSION

The PMM in 2030

Let's go back to Sarah.

You met her in the introduction—the product marketer who noticed her pipeline data going sideways because her buyers had started deploying AI agents. Website traffic up, form fills down. Prospects showing up to discovery calls knowing things they shouldn't have known. Enterprise deals closing in three weeks instead of eleven. Sarah's first instinct was that something was broken. What had actually broken was her mental model of how buying works.

Here's what happened next.

Sarah didn't panic. She didn't ignore it. She did what the best PMMs do: she got curious. She started testing. She ran her own product pages through an AI agent and asked it to evaluate her platform against a set of buyer requirements. The agent came back with a middling assessment—not because the product was middling but because the positioning was too vague for the agent to evaluate confidently. "The vendor appears to offer analytics capabilities but this is not explicitly confirmed." Sarah could hear her own marketing copy being graded, and the grade wasn't great.

She rewrote the product pages. Not in marketing-speak but in capability-speak—specific, structured, evidence-backed claims that an agent could parse and a human could trust. She ran the agent test again. The assessment improved dramatically. She ran it against three competitors. Two of them had the same vagueness problem she'd had; one—the competitor she'd always considered her biggest threat—had pages that were already agent-optimized. That told her something about which competitor was paying attention.

She built a CI monitoring system—nothing fancy at first, just a morning routine where she fed competitor news to Claude and asked for a structured analysis of what had changed. Within two weeks, she caught a competitor's quiet repositioning away from a market segment they'd been losing—a signal that would have taken months to surface through her old quarterly review process. She briefed her sales team that afternoon. They adjusted their messaging in two active deals. Both deals closed.

She redesigned her sales enablement model. Instead of producing static battlecards that decayed between quarterly updates, she built a knowledge base—her competitive intelligence, her product documentation, her customer case studies, her pricing data—and connected it to a RAG-powered system that sales reps could query in natural language. The first time a rep used it to prepare for a competitive call and came back saying "that was the best competitive brief I've ever had," Sarah knew she was

onto something.

Six months later, Sarah got promoted. Not to a bigger PMM role—to the newly created position of VP of Go-to-Market Intelligence, reporting to the CMO. Her charter: design and scale the agent-powered go-to-market infrastructure across the entire marketing organization. The promotion wasn't a reward for learning AI tools. It was a recognition that Sarah had understood, before almost anyone else at her company, that the PMM's job was no longer to produce deliverables. It was to build systems that produce better deliverables, faster, grounded in better intelligence, directed by better judgment.

• • •

Sarah's story is composited—drawn from real PMMs we know, real transformations we've witnessed, compressed into a narrative arc for clarity. But every element of it is happening right now, in 2026, at companies large and small. The PMMs who are moving fastest are already operating at a level that their peers from two years ago wouldn't recognize.

The Org Design Shift

The hardest part of the six percent problem is not the diagnostic — it is what the diagnostic implies. If AI integration is real, it does not produce a PMM organization that is the same size doing more work. It produces a PMM organization that is structured differently, staffed differently, and optimized for a fundamentally different mix of activities. The org design shift is not optional; it is the logical endpoint of genuine integration.

Figure 4: The Org Design Shift

BEFORE: Traditional PMM Org		AFTER: AI-Integrated PMM Org	
60–70 %	Commodity execution Content, research, templates, first drafts	~5%	AI-supervised commodity Agents produce; PMMs review and direct
15–20 %	Enablement ops Battlecard maintenance, training, RFPs	20–25 %	Competitive intelligence Living systems, agent-monitored, always-on

Figure 1: The Org Design Shift

So what does the PMM look like in 2030? We think the answer is: **bigger**. Not bigger in headcount—though we don't think headcount shrinks as much as the doomsayers predict—but bigger in scope, bigger in strategic influence, bigger in the breadth of the business they touch.

The PMM of 2030 is part **strategist**—making the judgment calls about positioning, competitive response, pricing, and market entry that determine whether a product wins or loses. Part **agent architect**—designing and maintaining the AI systems that produce the intelligence, the content, the enablement, and the launch operations that used to require entire teams. Part **customer psychologist**—understanding buyers at a depth that no agent can replicate, because understanding requires

having been in the room, read the body language, and heard what wasn't said. And part **organizational politician**—navigating the stakeholders, building the coalitions, and creating the alignment that turns good strategy into executed strategy.

That's a more demanding job than the PMM job of 2023. It requires a wider range of skills, a higher level of judgment, and a comfort with tools and systems that most PMMs didn't grow up with. But it's also a more rewarding job—more creative, more strategic, more directly connected to business outcomes. The mechanical work that used to consume 60% of the week is handled by machines, and the human gets to do the parts of the job that are actually interesting.

• • •

Chris started writing about data-driven marketing twenty years ago, and every book he's written has been about the same question: what happens to the people when the technology changes? In *Data Driven*, the answer was that the people who learned to use data platforms built careers that lasted decades. In *Customer Data Platforms*, the answer was that the people who understood identity and personalization became the most valuable members of their marketing organizations. In *Agent Driven*, the answer is that the enterprises that embrace agentic AI will outperform the ones that don't.

And in this book, the answer is personal. It's about you. Your career, your skills, your daily workflow, your Monday morning.

The agentic era is not something that's happening to product marketing. It's something that product marketers get to shape. The frameworks in this book—the three clusters, the dual-audience positioning model, the intelligence-layer enablement approach, the T-shaped PMM—are starting points, not endpoints. They're based on what we're seeing in early 2026, and the landscape will continue to evolve faster than any book can keep pace with. What won't change is the fundamental dynamic: machines are getting better at execution, which means humans need to get better at judgment. The PMM who invests in both—who builds the agent fluency and deepens the human skills—is the PMM who thrives.

We'll end with something we believe to be true, even though we can't prove it yet: the best product marketers five years from now will be more influential, more strategic, and more essential to their companies than the best product marketers today. Not despite AI. Because of it. The machines take the drudge work. The humans keep the interesting problems. And the interesting problems—the judgment, the empathy, the narrative craft, the organizational navigation—are the problems that make a career worth building.

That's the future of product marketing. Go build it.



THE CMO PERSPECTIVE

If there is one piece of advice we'd leave with every PMM who has read this far, it's this: don't wait. Don't wait for your company to give you AI tools—get them yourself. Don't wait for someone to redesign your workflow—redesign it yourself. Don't wait for the perfect agent platform to emerge—build something imperfect with what exists today. Don't wait for permission to experiment—experiment, and let the results earn you the permission.

The PMMs who are going to thrive in the agentic era are the ones who started before they were ready. They built a janky CI monitoring system out of RSS feeds and Claude prompts and showed it to their sales team before it was polished. They ran the agent test on their own product pages and had the uncomfortable realization that their positioning was invisible to AI buyers. They stopped producing static battlecards and started building knowledge bases, even when nobody asked them to. If you're waiting for certainty, you'll wait too long. The PMMs who move first—imperfectly, iteratively, with the willingness to be wrong and adjust—are the ones who will define what this discipline becomes. Go build it.

KEY TAKEAWAYS

- Start now, start imperfect: the PMMs who experiment first are already operating at a level their peers are still planning for.
- Build systems, not artifacts: knowledge bases over battlecards, continuous intelligence over quarterly reviews.
- The 10x PMM isn't working ten times faster—they're operating in a fundamentally different category of strategic impact.
- The discipline is being rewritten. The practitioners who write it will be the ones who showed up early and built.