

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.