PART II: DECONSTRUCTING THE PMM JOB

CHAPTER THREE

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.*

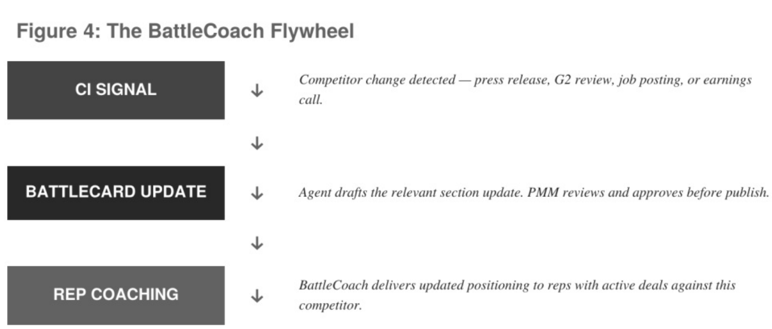

Figure 1: The BattleCoach Flywheel

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 *No system*	STAGE 2 PERIODIC *Quarterly cycle*	STAGE 3 MONITORED *Alerts + triggers*	STAGE 4 LIVING SYSTEM *Always-on agents*
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.

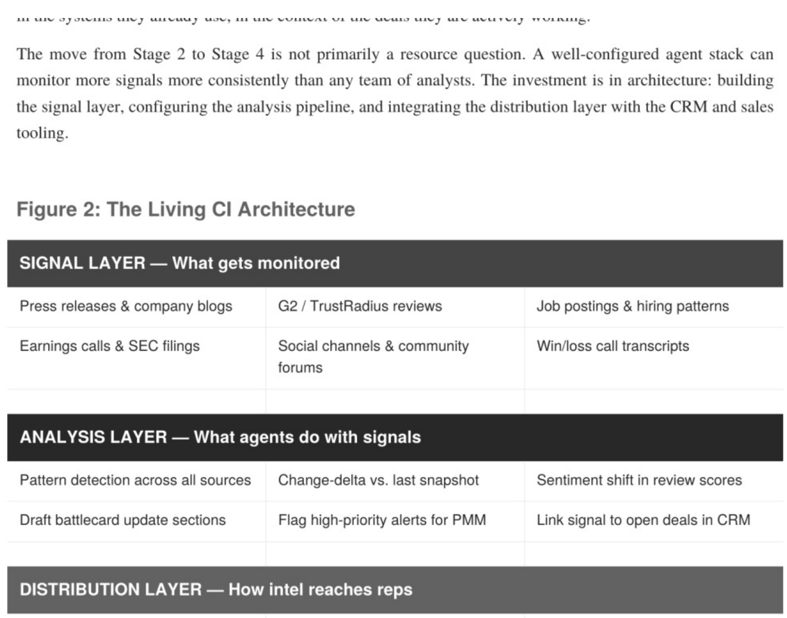

The move from Stage 2 to Stage 4 is not primarily a resource question. A well-configured agent stack can monitor more signals more consistently than any team of analysts. The investment is in architecture: building the signal layer, configuring the analysis pipeline, and integrating the distribution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