CHAPTER ELEVEN

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.

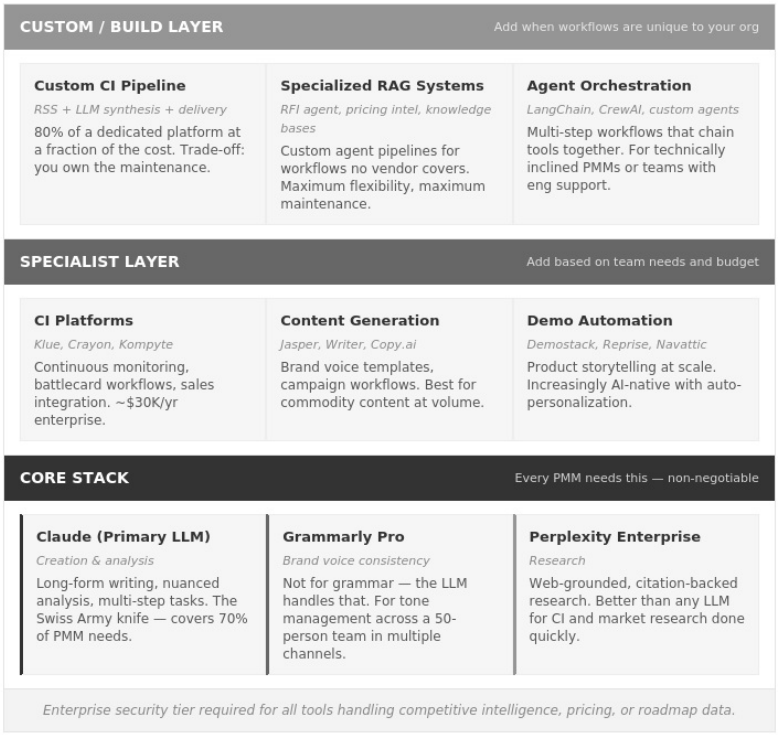

> **The principle:** Start with the general-purpose LLM for everything. Learn which workflows are worth investing in. Then specialize. You can always add layers — it's harder to untangle a Frankenstein of point solutions.

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.

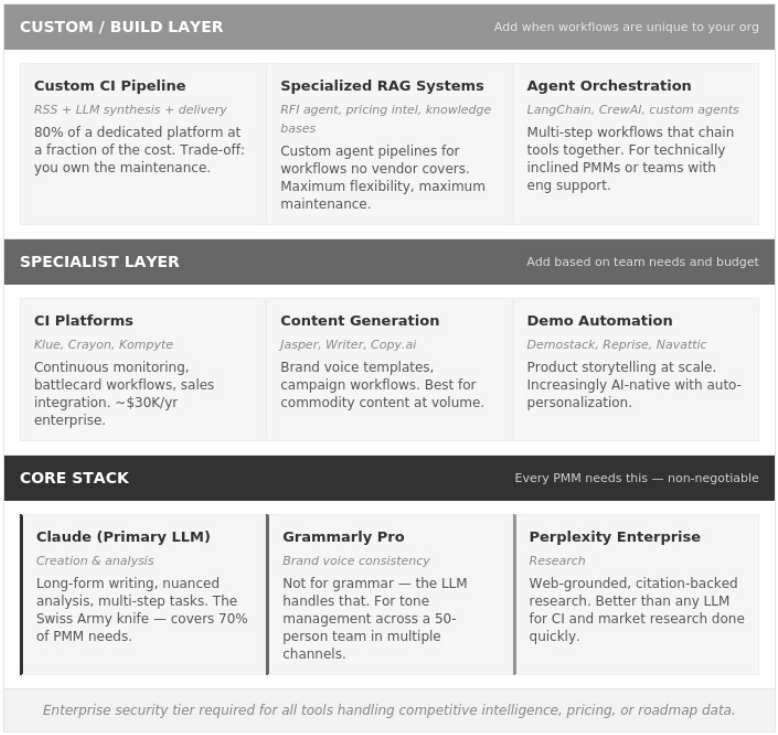

> **The principle:** Start with the general-purpose LLM for everything. Learn which workflows are worth investing in. Then specialize. You can always add layers — it's harder to untangle a Frankenstein of point solutions.

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.

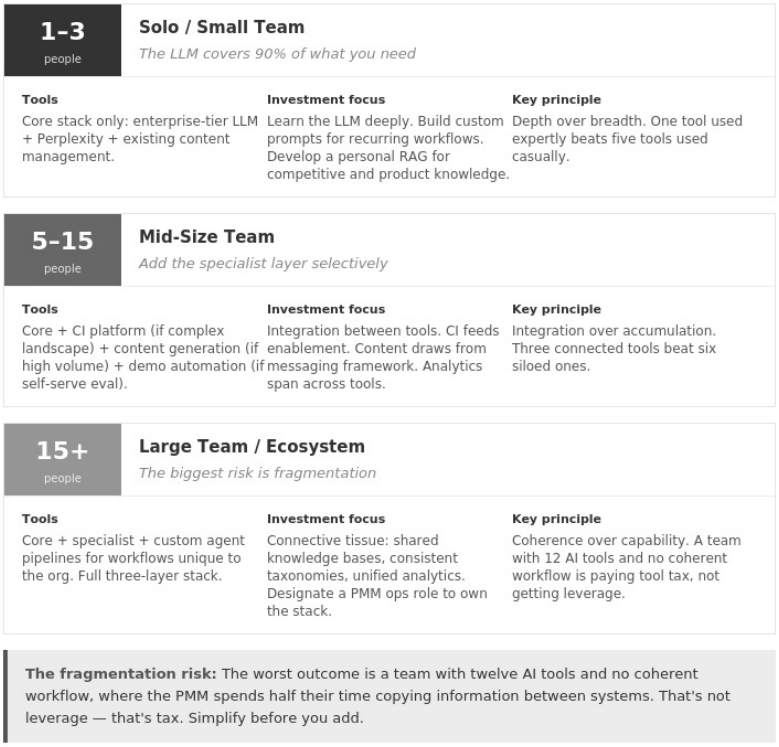

Figure 3. The Stack by Team Size. The core is universal. The specialist layer scales with need. The custom layer adds flexibility at maintenance cost. At every scale, coherence beats capability.

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.