The PMM Tech Stack:
A Practitioner's Evaluation

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

Over 200 AI tools in "AI for marketing" on G2 alone. We put together an acquisition plan for 53 people at SAP — it was supposed to take a week, it took three. This chapter cuts through the noise with a practitioner's framework for building a coherent stack.

> *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, bottom-up. Start at the core. Add upward only when you've exhausted the layer below.

CUSTOM / BUILD LAYER		
Custom CI Pipeline RSS+LLM synthesis	**RAG Systems** RFI agent, knowledge bases	**Agent Orchestration** LangChain, CrewAI
SPECIALIST LAYER		
CI Platforms Klue, Crayon, Kompyte	**Content Gen** Jasper, Writer, Copy.ai	**Demo Automation** Demostack, Reprise, Navattic
CORE STACK		
Claude (LLM) Creation & analysis — covers 70%	**Grammarly Pro** Brand voice across 50-person team	**Perplexity Enterprise** Citation-backed research

Figure 1. The core covers 70% of PMM needs. Specialist adds category capability. Custom adds flexibility at maintenance cost.

Figure 2: The Build vs. Buy Decision

Deciding factor: maintenance burden, not capability.

Workflow	Rec	Rationale
CI monitoring	BUY	Stable workflow. Platform handles RSS, parsing, alerting.
Content production	BUY volume / LLM strategic	Templates for commodity. LLM for nuance and strategic work.
Demo automation	BUY	Different class of problem. Increasingly AI-native.
CI synthesis	BUILD	Experimental, unique. Custom pipeline = exact synthesis you need.
RFI knowledge base	BUILD	Your responses + your docs = unique corpus no vendor has.
Everything else	START WITH LLM	Fastest way to learn which workflows are worth investing in.

Figure 2. Buy for stable workflows. Build for unique workflows. Start with the LLM for everything else.

Figure 3: The Stack by Team Size

What to deploy at each scale. Core is universal. Integration challenges multiply with team size.

01 Solo / Small (1–3)
LLM covers 90%

Core stack only. Learn the LLM deeply — custom prompts, context window management, personal RAG. Depth over breadth: one tool used expertly beats five used casually.

02 Mid-Size (5–15)
Add specialists selectively

Core + CI platform + content gen + demo. Focus on integration: CI feeds enablement, content draws from messaging framework, analytics span tools. Integration over accumulation.

03 Large Team (15+)
Biggest risk is fragmentation

Full three-layer stack + PMM ops role to own it. Connective tissue: shared knowledge bases, consistent taxonomies, unified analytics. Coherence over capability.

Figure 3. At every scale, coherence beats capability. Twelve tools with no workflow = tool tax.

Figure 4: The Stack Evaluation Framework

Three criteria plus the security conversation.

01 **Leverage**
Output per person

Does one PMM produce what two required, or the same output 50% faster? If the tool adds a step without removing one, it's net negative.

02 **Quality**
Maintained or improved

Does output quality hold or improve? A tool that trades quality for speed is a bad investment — downstream cost of bad intel exceeds time savings.

03 **Coherence**
Workflow integration

Does it fit existing workflow or create a silo? A PMM spending half their time copying data between tools is paying tool tax, not getting leverage.

Security: Enterprise tier required. Consumer-tier LLMs typically include data in model training. Enterprise tiers (Claude Team, ChatGPT Enterprise) include contractual guarantees. For PMM teams handling CI and pricing, enterprise tier isn't optional. Go to IT proactively with a proposal.

Figure 4. Leverage, quality, coherence + enterprise security. The strongest pitch is a before/after from your own workflow.

Key Takeaways

- The core stack is three tools: Claude, Grammarly Pro, Perplexity Enterprise. Covers 70% of PMM needs.
- 200+ tools, zero clarity. A coherent stack beats a Frankenstein of point solutions.
- Buy for stable workflows, build for unique workflows, start with the LLM for everything else.
- Scale by team size: depth (solo), integration (mid), coherence (large). PMM ops owns the stack at 15+.
- Evaluate on leverage, quality, and coherence. A tool that trades quality for speed is a bad investment.
- Enterprise security tier is non-negotiable. Go to IT proactively with a responsible business case.

The CMO Perspective

The tools conversation is the one PMM teams bring to leadership most frequently. Three criteria matter: leverage, quality, coherence.

- **Leverage** — does one person produce what two previously required?
- **Quality** — maintained or improved? Faster at lower quality is a bad investment.
- **Coherence** — fits the workflow or creates another silo?
- **Evidence over pitch** — the strongest signal isn't the vendor deck, it's a before/after from your own workflow.

> ***A small core stack plus selective specialist tools*** *beats a dozen disconnected point solutions.*