How to Build a GTM Enablement AI Project: A Blueprint for Product Marketers

A field guide to how one product marketing team set up an AI workspace to give Sales, Customer Success, and PMM a single place for verified answers, deal support, and on-brand content. Company, product, and system names are replaced with generic terms.

The Big Picture

What the project does

The project works as an **internal GTM enablement engine**. Reps, AEs, SDRs, CSMs, and PMMs ask questions or request drafts in plain language. The project then:

1. **Turns the request into a precise question** about the company's products, data, customers, or GTM motions, even when the user doesn't name the company.
2. **Sends it to the right specialist skill** (coverage, methodology, pricing, RFPs, prospecting, content, and so on).
3. **Answers only from approved internal sources**, following a fixed source-of-truth order.
4. **Applies standing guardrails** on claims, scope, confidentiality, and approval before anything is returned.
5. **Returns output in a set format**: a short verified answer, or a draft labeled for review.

The core idea

> One router, many narrow specialists, one shared rulebook.

The project instructions act as a **router and policy layer**. Each **skill** does one job well and names its own sources. Guardrails and voice rules apply to every skill.

Why build it this way

- **Trust:** Every answer traces back to an approved source. If no source supports an answer, the system says "unknown from company sources" and doesn't guess.
- **Consistency:** Every rep gets the same positioning, claim limits, and formatting.
- **Speed:** Routine questions (coverage, methodology, finding a deck) get answered without a Slack thread to PMM.
- **Scale:** PMM knowledge becomes reusable skills instead of one-off answers.

Architecture: Five Layers

Layer	What it is	What it controls
1. Project instructions	The system prompt for the workspace	Role, scope, source order, routing, guardrails, voice, output format
2. Skills	Modular instruction packs, each with a defined trigger and its own sources	How one type of request gets handled from start to finish
3. Connectors	Live links to internal systems	Product-truth database, CRM, wiki, shared drive, chat, email, ticketing
4. Knowledge files and wiki	Project files plus a shared memory of decisions and context	Extra detail and continuity across sessions. Not the main source
5. Scheduled refreshes	Recurring jobs that re-sync skills with their source docs	Stop skills from drifting when positioning or coverage changes

Layer 1: Project Instructions (the Rulebook)

Role and scope

- Name the audience: Sales, AEs, SDRs, CS, PMM.
- Assume every request is about your company, even when the company isn't named.
- Set a **hard scope boundary**. For example, a sister business unit's products are out of scope. When asked about them, say so and stop.

Source of truth, in priority order

1. **The matching skill's own named sources.** If a skill defines its source of truth (for example, "methodology answers come only from these five wiki pages"), that rule wins.
2. **The internal product-truth database** (product definitions, coverage, sellability) when no skill covers the question.
3. **Systems of record for account and deal context:** CRM, chat, email, shared drive, wiki, ticketing.
4. **Project files and the knowledge wiki**, for extra detail only.

If none of these supports an answer, reply "unknown from company sources" and stop.

Web-use policy

- **Never** use the public web for your own product facts: coverage, methodology, pricing, roadmap, or capabilities.
- **Allow** the public web only for prospect and account research (company news, priorities, people). Label it as external context and keep it apart from internal facts.

Routing table

Request type	Routed to
Coverage, category, retailer, "do we have data on X"	Coverage availability checker
Methodology, "how do we measure X"	Methodology Q&A
Pricing, quotes, renewals, discounts, churn risk	Renewals and commercial skill (internal only, needs sign-off)
Meeting or call prep	Meeting prep brief
Sequences, outreach emails, LinkedIn, call scripts	Prospecting
Customer questions, objections, "why us"	Customer call answer coach, grounded in current core positioning
RFPs, RFIs, questionnaires	RFP response assistant
Custom research	Custom research skill
Advisory product questions or content reviews	Advisory product FAQ and content reviewer
Finding decks, one-sheets, guides; survey sample sizes	Enablement search
One-pagers, letters, docs	One-page content builder
Any branded file, deck, PDF, or sheet	Design formatter
Raw data, exports, API, SFTP	Data delivery and integrations
Lifecycle and pipeline stages (feasibility, implementation, offboarding)	Pipeline stages skill

Route silently. The answer never names skills, tools, or internal calls. If a skill or connector fails or returns partial data, the answer says so under **Caveats** and lowers its confidence.

Standing guardrails

- **State country scope.** Coverage in one country doesn't prove coverage in another.
- **Separate "tracked" from "ready to sell."** Anything without a published coverage page needs Product confirmation before it's promised to a customer.
- **Name each product's real nature.** A modeled estimate is not a point-of-sale feed. Before any comparison, align category, retailer, SKU, and time period.
- **Treat roadmap items as directional.** Product must confirm them before customer use.
- **Protect confidentiality.** Don't name customers in external copy without authorization. Keep fees and deal terms internal.
- **Keep client-facing documents clean.** Don't expose internal routing, codes, or vendor details.
- **Draft only.** Never send, post, or update CRM records without explicit approval.
- **Review before release.** All customer-facing output is a draft for internal review.

Voice

- **Internal answers:** blunt, plain, short sentences. No jargon, hype, or praise. Lead with the biggest risk, gap, or constraint.
- **Customer-facing copy:** formal, warm, benefit-led, and aligned to positioning and brand voice.
- **Reinterpreted asks:** open with "Heads up -" and state the assumption.
- **Vague asks:** open with "Heads up - ambiguous request" and ask for the category, brand or account, and country.

Output format

Quick answers (coverage, methodology, account lookups):

- A bold title naming the topic.
- A one-line header: source, date checked, market, and **Confidence: high / moderate / low / unknown**.
- For coverage, a compact `Product Line | Status` table first, then no more than 2-3 sentences.
- Must fit on one screen. Add **Caveats** only when needed.

Drafts and deliverables (emails, sequences, decks, docs, briefs):

- Follow the skill's own format.
- Open with a one-line summary of what was built and what needs review.
- Link every internal source used.

Layer 2: The Skill Library

Each skill has a **trigger** (when it loads), **named sources** (where it may look), a **workflow** (the steps it follows), and **boundaries** (what it won't do). The skills fall into five families.

A. Verified Answers (product truth)

Skill	Trigger	What it does
Coverage availability checker	A category is named with "what's available," or a question about roadmap, past releases, retailer coverage, or a specific pricing dataset	Checks coverage master lists on the wiki against the product-truth database. Returns a status table by product line and country, plus collection start dates and cadence. Answers roadmap and release-history questions from the roadmap tracker and release notes
Methodology Q&A	The word "methodology," or "how do we measure X"	Answers **only** from a fixed set of methodology overview pages, with an inline citation to the specific page. No web and no other docs. If those pages don't cover it, says so and stops

Skill	Trigger	What it does
Product coverage and definitions	"What do we sell or track," product definitions, ready-to-sell status	Explains coverage, cadence, and status, and sends issues to the right team (Product, Data Ops, Engineering, Data Science, Support)
Capabilities fit router	"Can we do X?"	Labels the request as standard, custom, beta, POC, unsupported, source-blocked, or a defect, and names who owns it
Answer calibration	Any question about confidence or readiness	Scores source strength, separates fact from inference, and marks the answer ready for internal, customer, or leadership use, or for owner confirmation
Product methodology explainer	Deeper methodology or analytics concepts (elasticity, demand curves, refresh cadence)	Explains them in customer-safe terms

B. Customer Conversations

Skill	What it does
Customer call answer coach	Writes confident, customer-safe answers to objections, methodology questions, coverage questions, data-quality concerns, and timeline questions, grounded in the current core positioning doc
Customer operations response	Writes sanitized external explanations of operational issues (data quality, delivery, dashboard bugs) without exposing internal runbooks
Internal status comms	Writes internal updates, customer-safe status notes, "don't go dark" check-ins, and exec summaries
Helpdesk and runbook triage	Picks the right runbook, validates the issue on public surfaces, and drafts evidence-backed tickets

C. Deal Support

Skill	What it does
Meeting prep brief	Takes a company, website, attendees, meeting type, and prior knowledge. Returns a **five-minute brief mapped to capabilities**, 12 months of CRM deal history, and **three tailored email hooks**. Web research is labeled as external context
Prospecting sequence builder	Builds an **8-step multi-channel sequence** (email, LinkedIn, phone) with persona mapping, jobs-to-be-done framing, hook selection, a use-case template, competitor battlecards, and voice rules
RFP response assistant	Uses a **3-level fallback, always in order**: (1) methodology and coverage skills, (2) the product-truth database, (3) the RFP answer library, used only for narrative framing on top of verified facts. Flags pricing, contract, security, and roadmap answers as unverified until signed off. Delivers inline or into a spreadsheet. **Runs a positioning guardrail pass before returning anything.** Never submits externally
Renewals, pricing, commercial	Handles renewal context, churn risk, concessions, and handoffs. Pricing stays internal and needs sign-off
Pipeline stages	Explains lifecycle stages (feasibility, implementation, offboarding) and which team owns each piece of work

Skill	What it does
Data delivery and integrations	Says whether a requested export, API, SFTP, or file shape is standard, custom, sold, or out of scope

D. Content and Brand

Skill	What it does
Core positioning guardrail	**Loads first** whenever copy is involved. Checks any draft (deck, email, FAQ, RFP answer, battlecard, AI output) against core positioning docs for every product line. A scheduled job keeps it in sync with those docs
One-page content builder	Builds external docs (one-pagers, briefs, FAQs, letters) on a **single branded letterhead template**. Applies brand fonts, colors, and approved logos, plus **accessible tables** (native header rows, color contrast that passes WCAG, consistent alignment). Restricted to one output format. Always flags the result for internal review
Design formatter	Applies the full brand system to any deck, doc, sheet, PDF, or image
Transcript-to-use-case	A **three-stage pipeline**: (1) pull grounded evidence from a call transcript and rank candidate use cases, (2) write battlecard slide copy and speaker notes, (3) run a senior PMM positioning review
Advisory product content reviewer	**Review only, never writes new copy.** Two personas: a senior content marketer scores a red/yellow/green guardrail matrix and tiered findings, then **pauses**. A senior PMM then validates each finding against four canonical sources, explains why, and handles pushback by citing sources
Advisory product FAQ	Answers only what the four canonical docs state, with a hyperlinked citation on every answer. No inference
Custom research	Answers questions about custom surveys, focus groups, and mystery shopping, and points to decks, sample proposals, and SOWs

E. Discovery and Analyst Support

Skill	What it does
Enablement search	A structured intake (product line, then resource type, then optional keywords). Searches only the approved drive folders and returns a ranked list of direct links. Also answers survey sample-size questions by product, category, and quarter
Line-review data sheet builder	Fills a multi-tab analyst workbook from the product data, with tab QA and insight notes

Patterns Worth Copying

1. **Give every skill a single, named source of truth.** The narrower the source list, the more trustworthy the answer. "Only these five pages" beats "search everything."
2. **Use fallback ladders, not open search.** For example: verified product facts, then the product database, then the answer library for framing only.
3. **Build in stop rules.** "If not in the source, say so and stop" is the most important line in many skills.

4. **Split creation from review.** Use separate skills, or separate personas inside one skill, with a pause between finding issues and rewriting.
5. **Make the guardrail the last step.** Any skill that produces external copy hands its output to the positioning guardrail before returning it.
6. **Show confidence to the reader.** Confidence level, date checked, and market scope appear in every quick answer.
7. **Route silently.** Users see answers, not plumbing.
8. **Default to drafts.** Nothing is sent, posted, or written to the CRM without explicit approval.
9. **Layer skill scopes on purpose.** Skills can exist at personal, project, and organization level. Decide which version wins (personal, then project, then org, then built-in) and retire duplicates so they don't drift apart.
10. **Schedule source syncs.** Positioning and coverage change. A recurring job that compares each skill with its source docs keeps answers current.
11. **Keep skill changes in draft until the owner approves.** Treat skills like product copy: version them and review before publishing.

How to Build Your Own: Step by Step

1. **List the questions your field team asks most.** Group them into families: product truth, customer conversations, deal support, content, discovery.
2. **Name the source of truth for each family.** If one doesn't exist, fix that first. The AI can only be as good as the docs it reads.
3. **Write the project instructions:** role, scope, source order, web policy, routing table, guardrails, voice, output formats.
4. **Build the guardrail skill first.** Everything else that produces copy depends on it.
5. **Build the highest-volume answer skills next.** Coverage and methodology usually remove the most PMM interruptions.
6. **Add deal-support skills** (meeting prep, prospecting, RFP), each with its own workflow and fallback ladder.
7. **Add content skills** with template, brand, and accessibility rules built in.
8. **Connect the systems:** product database, CRM, wiki, shared drive, chat.
9. **Pressure-test** with real rep questions. Check sources, stop behavior, confidence labels, and that no internal details leak into customer-facing drafts.
10. **Schedule refreshes and set an owner** for each skill.

Starter Template: Skill Definition

```
Name:            <short, specific>
Trigger:         <exact phrases or conditions that load this skill>
Not for:         <adjacent requests that belong to another skill>
Sources:         <named docs/pages/systems, in priority order>
Workflow:        <numbered steps, including any pause or review stage>
Stop rule:       <what to say when sources don't support an answer>
Output format:   <table / brief / draft / workbook, plus required header fields>
Guardrails:      <claim limits, confidentiality, approval gates>
Handoff:         <which skill runs next, e.g., positioning guardrail>
Owner + refresh: <who maintains it and how often it re-syncs>
```