




TWEGS RESEARCH REPORT • JULY 2026

Product Marketing in the Age of AI

How AI-mediated buying, cheap production, and the AI product wave are rewiring the PMM discipline. With playbooks by industry and by segment.

94%

of B2B buyers used AI in their most recent purchase

2.5

vendors on the average shortlist, down from 3.2

51%

start research in an AI chatbot, not a search engine

A note on method: figures from vendor-published research should be treated as directional. Where several independent studies converge, confidence is higher than for any single number. Statistics that could not be cross-checked were left out or tied to their one source. The full source list is at the end.

JULY 2026 • FIRST EDITION • [TWEGS.COM/PRODUCT-MARKETING-IN-THE-AGE-OF-AI](https://twegs.com/product-marketing-in-the-age-of-ai)


CONTENTS

Executive summary	3
Part I: The structural shift	4
1. The new buying reality	5
2. The AI visibility imperative: from SEO to AEO	6
3. The PMM role, redefined	8
Part II: The craft, rebuilt	10
4. Market and customer research	11
5. Positioning and messaging	11
6. Launch and go-to-market execution	13
7. Sales enablement and competitive intelligence	13
8. Pricing, packaging, and monetization	14
9. Measurement, attribution, and proof of impact	14
Part III: Industry and segment playbooks	16
10. The segmentation framework	17
11. Financial services and fintech	18
12. Healthcare and life sciences	19
13. Cybersecurity	20
14. Developer tools and infrastructure	20
15. Horizontal B2B SaaS	21
16 to 19. The remaining verticals, in brief	22
Part IV: The operating model	24
20. Team design, skills, and tooling	25
21. Governance, risk, and trust	25
22. A twelve-month roadmap	26
23. Outlook: 2027 and beyond	27
Sources	28
Appendix: key statistics	29

Exhibits 1 to 10 appear alongside the chapters they support. Page numbers refer to this print edition.

Executive summary

Product marketing is going through its largest structural change since the discipline emerged. Three shifts arrived at once, and any one of them would have reshaped the job on its own.



- 01 The buyer moved into the machines.** Forrester's 2026 Buyers' Journey Survey of roughly 18,000 business buyers found that 94% used AI during their most recent purchase, and AI answer engines now outrank vendor websites, sales reps, and product experts as the number one research source. 69% chose a different vendor than they initially planned because of AI guidance, and a third bought from a vendor they had never heard of before an AI put it on the shortlist. The first impression a company makes is increasingly an AI-generated answer to a prompt the vendor never saw, which moves the decisive work upstream of the funnel.
- 02 Production collapsed in cost, so the value moved to judgment.** Content, battlecards, launch collateral, competitive teardowns, research synthesis: the artifacts that used to consume most of a PMM team's week now take hours instead of days. This does not make product marketing less valuable. It concentrates the value in what to say, to whom, why it is true, and whether the claim survives contact with a procurement team, a regulator, or an AI agent fact-checking it against public documentation.
- 03 The product being marketed is now AI itself.** AI features, AI agents, and AI-native platforms are sold into a market saturated with identical claims. Positioning "AI-powered" is now positioning nothing. Buyers scrutinize AI claims harder than any other capability, and in regulated industries those claims carry enforcement risk.

Part I of this report maps the structural change in buying behavior and discoverability. Part II rebuilds the core PMM craft around the new reality. Part III works through the picture industry by industry and segment by segment. Part IV covers the operating model: team design, skills, governance, and a twelve-month roadmap.

AI has not replaced product marketing. It has removed the places where mediocre product marketing could hide.

PART I

The structural shift

The buyer, the cost of production, and the product itself all moved at once. Chapters 1 to 3 map what changed and how fast.

94%

of buyers used AI in their most recent purchase

#1

AI answer engines as a research source, ahead of websites and reps

90%

of B2B buying agent-intermediated by 2028, per Gartner

Chapter 1. The new buying reality

The pre-AI B2B buying journey was already vendor-hostile. Buyers spent most of their journey away from sales teams and arrived with a shortlist largely formed. Gartner's time-allocation research put independent online research at 27% of total buying time against 5 to 6% spent with any single vendor, roughly a five-to-one ratio of self-directed research to direct vendor contact. What AI changed is what happens inside that independent research window.

The 2026 numbers are consistent across independent studies even where the exact figures differ:

Forrester (January 2026, n of about 18,000): 94% of buyers used AI in their most recent purchase. 55% compare vendors inside AI tools, 54% research products there, and 47% build internal business cases before any vendor contact.

G2 (March 2026, n of 1,076): 51% of software buyers start research in an AI chatbot. 71% rely on AI chatbots for software research, up from 60%. 83% feel more confident in their final choice, and four in five say AI accelerated their decision.

Multi-source analysis (Loganix and Averi, March 2026, 680 million citations): 73% of B2B buyers use AI tools in purchase research. AI search traffic converts at 14.2% versus 2.8% for Google organic, a 5.1x advantage.

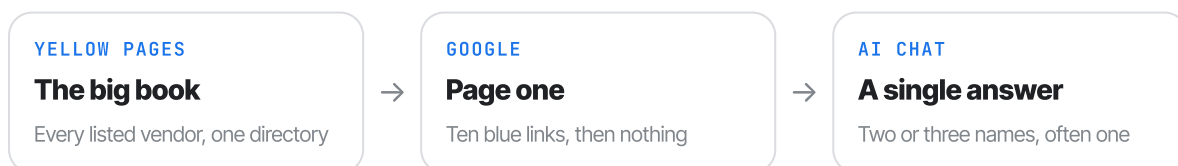
Marketing Graham's 2026 tech buyer study (n of 792): AI assistants scored highest among digital research sources at 2.31, ahead of search engines at 1.92, while peer recommendations remained the most trusted source overall at 2.88.

Shortlists are compressing. Apollo's 2026 data puts the average vendor shortlist at roughly 2.5, down from 3.2 a few years earlier. G2's Tim Sanders frames it as the third compression era: the Yellow Pages compressed the market into the big book, Google compressed it into the first page, and AI chatbots are compressing it into a single answer (Exhibit 1). When the answer surfaces two or three names, being fourth is functionally being invisible.

EXHIBIT 1

Three compressions of the vendor market

How many vendors the buyer ever sees, by era



Framing: Tim Sanders, G2. Each era shrank the visible market; the current one shrinks it to the contents of one answer.

Discovery moved off vendor property. B2B companies report traffic declines of 10 to 40% as research migrates into AI answer engines. The shortlist is now assembled inside a system that operates without the vendor's website, forms, or retargeting. Traditional attribution captures only around 27% of the buyer journey. The rest happens in what practitioners call the dark funnel: AI conversations, private communities, peer discussions, review platforms.


Discovery and validation split into different games. Buyers use AI to discover and compare, then use peers, reviews, and demos to validate. A third of buyers purchased from a vendor they had never previously heard of, which is pure AI-driven discovery. But peer recommendations still score highest on trust, and 20% of buyers say AI made them less confident because of unreliable information. Winning requires being present in the AI conversation and credible in the human one.

The shift is uneven across buyers, which matters for segment strategy. Leadscale's 2026 analysis found 85% of buyers aged 25 to 34 use AI for supplier research versus 23% of buyers aged 55 to 64, a 62-point gap (Exhibit 2). Since the 25 to 44 cohort now holds most purchasing authority in mid-market companies, LLM-first research is becoming the default fastest in exactly the segment where deals are big enough to matter and buying processes light enough to move quickly.

EXHIBIT 2

AI use in supplier research, by buyer age

Share of buyers using AI for supplier research, 2026



A 62-point generational gap. The younger cohort holds most mid-market purchasing authority, so LLM-first research becomes the default where deals move fastest. Source: Leadscale, 2026.

Meanwhile the buying committee keeps growing. Forrester puts the average B2B purchase at 13 internal stakeholders plus 9 external participants. AI has not shrunk the committee. It has armed every member of it with the ability to run their own research, build their own comparison, and draft their own objections before the vendor knows the deal exists. Gartner projects that by 2028, 90% of B2B buying will be agent-intermediated in some form.

What this means for product marketing: the window where marketing can influence a purchase is earlier and narrower than most GTM strategies assume. 6sense's research shows the winning vendor was already on the buyer's day-one shortlist 95% of the time, and 80% of buyers contact their pre-existing favorite first. If shortlists form inside AI conversations fed by public content, third-party citations, review data, and analyst coverage, the strategic center of product marketing moves upstream: from converting demand to shaping what the machines and the market believe before demand ever surfaces.

Chapter 2. The AI visibility imperative: from SEO to AEO

Search engine optimization asked: how do we rank? Answer engine optimization (AEO) and generative engine optimization (GEO) ask: how do we get cited and recommended? These are different problems with different levers. The research base is young but converging on a few findings.

Authority dominates. Ahrefs' analysis of ChatGPT citation behavior found 65.3% of top-cited pages come from domains with DR80 or higher. Authority is built through earned media over time, which means PR, analyst relations, and third-party coverage now do double duty as machine-visibility infrastructure.

Statistics and citations lift inclusion. The Princeton and Georgia Tech GEO research (Aggarwal et al., SIGKDD 2024) found adding statistics improves AI citation rates by 30 to 40%, and citing credible sources improves citation probability further. Concrete, sourced, structured claims outperform adjectives.


Mentions beat backlinks. Brand web mentions correlate roughly three times more strongly with AI citation rates than backlinks (Spearman 0.664 versus about 0.2). The unit of AI-era visibility is the credible mention in a source the models trust, not the link.

AI traffic converts disproportionately. AI-referred visitors convert at roughly five times the rate of traditional organic and spend 68% more time on site. The implied endorsement of an AI recommendation carries weight: 85% of buyers think more highly of a vendor when an AI chatbot names it.

EXHIBIT 3

Conversion rate by traffic source

Share of visits that convert, 2026



AI-referred visitors convert about 5.1x better and spend 68% more time on site. Source: Exposure Ninja and SE Ranking, 2026.

Google search itself changed shape. AI Overviews appear on nearly half of searches as of early 2026 and consume roughly half the mobile screen, with significant click-through declines on queries where summaries appear. The conclusion the analyst community keeps repeating: value accrues to being the source the model cites, not the page that ranks third.

Buyer behavior has outrun marketing operations by a wide margin (Exhibit 4). Only 22% of marketers currently track AI visibility. Only about 26% plan to develop content specifically for AI citations. 64% are unsure how to measure AI search success, even as 72% of marketing leaders expect AI to surpass SEO as the primary visibility channel within three years. This gap is an arbitrage window. Categories where no competitor has systematized AI visibility are winnable cheaply.


CASE: THE COST OF OPTING OUT

A mid-market SaaS company positioned itself as "human-written, AI-free" and skipped AEO entirely. Organic traffic fell 34% in six months while cited competitors gained. Recovery cost roughly \$45,000 and four months spent rewriting more than 60 posts.

EXHIBIT 4

Marketers on AI visibility: belief vs practice

Share of marketers, 2026



72% expect AI to become the primary visibility channel within three years, yet only 22% track it today.

Source: Yext, 2026.

AI visibility is often parked with SEO or content teams, but its core inputs are product marketing artifacts:

Machine-legible positioning. Product pages, comparison pages, and documentation structured so an agent can extract what the product does, for whom, against what alternatives, at what price, with what proof. Vague benefit language fails twice, with humans and with machines.

Claim integrity. AI engines synthesize from many sources. Inconsistent claims across the site, review profiles, and press coverage produce muddled or hedged AI answers. Message discipline is now a technical requirement, not just a brand one.

Third-party proof orchestration. Review volume and recency, analyst mentions, credible press, and case studies with named metrics are the raw material AI engines cite. PMM owns the proof-point supply chain.

Prompt-space intelligence. The new keyword research is prompt research: what questions do buyers in your ICP actually ask AI, and what does the AI currently answer? Monitoring share of voice in AI answers across the major engines is becoming a standing competitive-intelligence function.

So what. Citations are a supply chain. Give it an owner, a baseline, and a quarterly target, the way search rankings had one in 2015.

Chapter 3. The PMM role, redefined

By 2026, AI tools reliably handle the bulk of PMM production work: first-draft messaging documents, launch collateral, battlecards, objection-handling guides, win/loss synthesis, research summarization, competitive change tracking, and localization. Gartner's 2026 Priorities for Product Marketing Leaders calls AI "the defining force" reshaping PMM team structure and GTM strategy, and names AI fluency the single most important investment, with skill gaps now the number one obstacle to adoption.

Product Marketing Alliance's 2026 state-of-the-function research describes the same picture from the practitioner side: widespread individual adoption, genuine caution at the organizational level, and time freed from content production, competitive tracking, and research synthesis. Adweek's analysis of AI's effect on marketing organizations adds the uncomfortable middle chapter: before any headcount effect, the damage shows up as role confusion and eroding confidence, as organizations talk "AI-first" while still rewarding old signals of seniority.

The consensus across Gartner, PMA, and practitioner writing converges on what becomes more valuable:

Positioning judgment and taste. Knowing a technically correct positioning statement feels wrong. Knowing a battlecard will be ignored because it does not sound like how sellers actually talk. Pattern recognition built on reps, not tooling.

Pricing, packaging, and segmentation strategy. The commercial decisions AI can inform but not own.

Cross-functional decision rights. PMM value in 2026 is proportional to how early in decision-making it gets involved. Teams pulled in at execution ship. Teams that shape strategy win.

Claim governance. As output scales, someone must own what the company is allowed to say, especially where AI claims meet regulators and where AI agents generate customer-facing language.

Evidence over volume. AI makes output cheap, which tempts teams toward volume as a proxy for value. Leaders who redirect freed time into pricing, segmentation, and category strategy will have a very different year than those who do not.

The trap is treating PMM as a content function. If product marketing is primarily producing documents, AI genuinely does replace much of it. If it is making the judgment calls that connect product, market, and customer, and being accountable for whether those calls were right, AI is leverage. Organizations are quietly sorting PMM teams into those two categories right now, and the sorting shows up in budgets: leadership increasingly demands a direct line from PMM activity to revenue.

A useful test for any PMM task in 2026: does this require knowing something about our customers, our product, or our market that is not written down anywhere? If yes, it is durable human work, and worth writing down, because it becomes proprietary context that makes your AI tooling better than a competitor's. If no, it should already be automated, and the PMM's job is quality control and distribution.

So what. Reinvest the freed production hours in pricing, segmentation, and category strategy, or the argument for the headcount goes away with the output it used to justify.

PART II

The craft, rebuilt

Research, positioning, launch, enablement, pricing, and measurement, rebuilt for a buyer who reads with a machine. Chapters 4 to 9.

30 to 40%

citation lift from adding statistics to content

5.1x

conversion advantage for AI-referred traffic

27%

of the buyer journey visible to traditional attribution

Chapter 4. Market and customer research

Research was the first PMM workflow AI transformed, because it is mostly synthesis over large volumes of unstructured input, which is exactly what language models do well.

What changed. Win/loss interviews, sales call recordings, support tickets, review corpora, community threads, and analyst notes can now be continuously synthesized rather than sampled quarterly. AI-first teams work from a steady flow of clustered signals, a live map of objections, competitor mentions, and unmet needs, instead of a research report that is stale by the time it circulates. Synthetic-data and AI-augmented panel techniques have matured enough that major research firms report 94 to 95% accuracy against ground truth for well-scoped audience questions, making early concept and message testing dramatically cheaper.

What did not change. Synthesis is not insight. The models compress. The PMM decides what matters. Two disciplines separate strong AI-era research practice from lazy practice. First, read the outliers yourself, because AI summarization systematically sands off the weird quote that contains the actual positioning insight. Second, keep provenance: know which claims come from customers' mouths and which are the model's interpolation, because you will be asked to defend the insight in a pricing or roadmap fight.

New research surfaces. Two research objects did not exist three years ago. Prompt-space research: what your buyers ask AI about your category and what the models answer, including how they describe your product, which competitors they pair you with, and which sources they cite. And agent-behavior research: as procurement teams begin using AI agents to screen vendors, understanding what those agents extract from your public materials becomes a testable, improvable input.

The stack pattern. The common 2026 architecture is a pipeline, not a tool. Call recordings and CRM notes flow into a transcription and tagging layer, a model clusters themes weekly, and the PMM reviews a digest plus raw outliers. Teams that built this report a compounding effect: every launch, pricing decision, and battlecard starts from live evidence instead of memory.

Chapter 5. Positioning and messaging

The core positioning problem of 2026 is that everyone's messaging is converging. When most first drafts come from similar models prompted with similar category language, differentiation by wording dies. Three practical responses:

Position on evidence, not adjectives. Statistics, named customers, and specific mechanisms both persuade humans and lift AI citation rates by 30 to 40%. "Reduces false positives 43% for teams running Splunk" beats "AI-powered security insights" in every channel that matters.

Position against the AI-generated default. Ask the major models what they would recommend in your category and why. Their answer is the composite of your market's received wisdom, which is the thing your positioning must beat. If the model's summary of your category is indistinguishable from your homepage, your homepage is doing nothing.

Own a claim with a falsifiable edge. Claims that can be checked survive. Claims that cannot get averaged away. AI agents increasingly verify sales claims against public documentation, and buyers use LLMs specifically to bypass marketing language and compare specs. Positioning that your own documentation contradicts is now discovered before the first call.

The "AI-powered" problem. The 2025 to 2026 period produced a surge of AI-washing: AI labels on every feature without articulated value. Buyer response has been sharp. Technology buyers now scrutinize AI capabilities harder than any other claim area, nearly all expect vendors to demonstrate AI capabilities concretely, and roughly a third report caution about AI features specifically because of privacy or ethical concerns. A counter-trend of "human-made, AI-free" branding emerged, which works as brand contrast in some consumer categories but fails badly when extended to refusing machine-readable content.

The workable messaging frame for AI capabilities has three layers (Exhibit 5). Enterprise buyers buy the guardrails as much as the capability. Security and compliance risk is the single top concern in tech purchasing, at 38%.

EXHIBIT 5

The three-layer frame for AI capability messaging

Every AI claim carries all three, in this order

1 Outcome

What improves, by how much, for whom.

2 Mechanism

What the AI actually does, in one honest sentence.

3 Guardrails

Data handling, human oversight, failure modes.

Skipping the third layer moves the conversation to the security questionnaire, where marketing no longer controls the story.

Messaging for two audiences: humans and machines. Every messaging artifact now has a second reader. The practical checklist:

Structure content so claims are extractable: clear headers, direct answers near the top, specs and pricing in parsable form, comparison content that names competitors honestly.

Maintain consistency across every surface the models ingest: site, docs, review responses, press. Contradictions produce hedged AI answers.

Feed the third-party layer deliberately: reviews with substance, analyst briefings, earned media in high-authority publications, because 65% or more of citations come from high-authority domains you do not own.

Keep a canonical claims file, the approved set of numbers, customer names, and superlatives, as the single source both your writers and your AI tooling draft from.

Chapter 6. Launch and go-to-market execution

Launches compressed. The production timeline that once forced a six-week runway now fits in days, which means the constraint moved to alignment and judgment: is the story right, is sales bought in, is the claim defensible.

What good looks like in 2026: AI drafts the artifact tree from a single approved messaging source, and the PMM edits for specificity, inserting the beta metric, cutting the unsupportable claim, replacing generic benefits with the customer quote that nails it. AI also runs logic checks over the launch plan itself, catching dependency conflicts and sequencing errors like enablement scheduled after release, a genuinely valuable use because humans managing forty tasks miss exactly these.

Because content is cheap, launch differentiation shifted to what cannot be generated: proof (benchmarks, named customers, live demos), access (community, executive engagement, analyst validation), and experience (interactive tours, sandboxes, the product visibly doing the thing). Release velocity changed the cadence question too. With product teams shipping continuously, tiered launches plus a stream of always-on, machine-visible release notes replaced the monolithic launch. Changelogs and release notes are heavily ingested by AI engines, an underrated visibility surface.

Launching AI products specifically: the evidence bar is higher. Buyers have learned that AI demos exaggerate. The launches that convert lead with constrained, verifiable claims plus a path to hands-on validation, and address the governance questions in the launch materials rather than leaving them to the security questionnaire.

Chapter 7. Sales enablement and competitive intelligence

Enablement content was the highest-volume, most-structured PMM output, and is therefore the most automated by 2026. The differentiated work moved up a level.

From documents to systems. Instead of shipping a battlecard, leading teams ship an always-current competitive layer: a governed knowledge base sellers query mid-deal, fed by the monitoring pipeline and reviewed by PMM. The PMM's job is the curation and the judgment calls encoded in it. What we say, what we never say, where we genuinely lose.

From training events to deal support. Buyers arrive at first contact more informed and less patient with early-stage education. Two-thirds now prefer engaging sales only late in the journey, so enablement shifts toward late-stage, high-stakes moments: business-case support, security and compliance narratives, executive alignment.

Honesty as strategy. Because buyers verify claims with AI against public sources, battlecards built on stretched claims now actively damage deals. The 2026-grade battlecard concedes real weaknesses and arms sellers with the trade-off narrative instead.

Competitive intelligence went from periodic to continuous. The standard pattern is automated monitoring of competitor sites, changelogs, pricing pages, job postings, and review streams, with diff-detection before AI summarization, feeding a weekly human-reviewed digest and a living battlecard layer. Two additions are new. First, AI-answer share-of-voice tracking: systematically prompting the major engines with buyer-realistic questions and logging who gets recommended, with what framing, citing which sources. This is now as fundamental as tracking search rankings was

in 2015. Second, counter-positioning against the models' picture of competitors: if the AI engines describe a competitor more generously than reality, the fix is publishable evidence the engines can ingest, not internal grumbling.

Chapter 8. Pricing, packaging, and monetization

Pricing is where AI-era product marketing gets hardest, because AI broke the dominant SaaS pricing model in two directions at once.

The seat problem. Per-seat pricing assumes value scales with human users. AI features and agents decouple value from headcount. An agent that resolves 40% of support tickets reduces the seats you can bill. The market's response has been rapid diversification: usage-based pricing, hybrid seat-plus-consumption models, and outcome-based pricing (per resolved ticket, per qualified meeting, per document processed). Each has a known failure mode PMMs must design around. Pure usage pricing makes costs unpredictable and triggers procurement resistance. Outcome pricing invites definitional disputes. Hybrids add complexity that lengthens sales cycles.

The cost problem. Unlike traditional software, AI features carry real marginal cost, which reintroduces gross-margin discipline into packaging. Inference costs have fallen about 90% over three years, which is what made production AI viable for mid-market deployment, but heavy-usage customers can still invert unit economics. The standard packaging levers: usage tiers with soft caps, premium AI SKUs, credit systems, and fair-use policies with teeth.

Segment behavior. Enterprise buyers pay for predictability and governance: committed-spend contracts with usage bands, plus paid tiers for audit, controls, and private deployment. Mid-market wants transparent calculators and self-service expansion. SMB pricing gravitates to simple bundles where AI is included, not itemized, because SMB buyers reject metering complexity outright.

The PMM mandate: willingness-to-pay research for AI capabilities, value-metric selection, continuous competitive price-position monitoring, and the migration story, meaning how existing customers move to new models without a churn event. In 2026 budget reviews, pricing work is consistently the PMM investment leadership values most, because it is the shortest line from the function to revenue.

Chapter 9. Measurement, attribution, and proof of impact

The measurement problem got worse before it got better. Most of the buying journey now happens in channels that emit no signal, while the channels that do emit signal, site traffic and form fills, are declining as research migrates off vendor property. Teams still managing to MQL volume are optimizing a shrinking, unrepresentative slice of the journey.

The measurement stack serious teams converged on:

AI visibility metrics as a first-class KPI family. Share of recommendation in buyer-realistic prompts across major engines, citation frequency and source mix, sentiment and framing of AI descriptions, and referral traffic from AI surfaces, which converts around five times organic.

Self-reported attribution, taken seriously. A free-text "how did you hear about us?" on every conversion point, mined with AI, is the only instrument that sees the dark funnel. Teams consistently find AI tools and communities dramatically over-indexed versus what click attribution shows.

Correlation over click-paths. Brand search volume, direct traffic, review velocity, and share-of-voice trends against pipeline creation, accepting that deterministic multi-touch attribution now describes roughly 27% of reality.

PMM-specific commercial metrics. Win rate in competitive deals, deal velocity where enablement engaged, pricing realization, launch-sourced pipeline. The metrics that survive contact with a CFO.

The honest caveat every PMM should carry into the boardroom: precision is lower than the dashboards imply, and pretending otherwise is how budgets get misallocated. The defensible posture is triangulation, multiple imperfect instruments pointing the same direction, plus periodic incrementality tests where spend is deliberately varied.

PART III

Industry and segment playbooks

The same playbook fails differently in every vertical. Chapters 10 to 19 work segment by segment through what changed and what to do about it.

94% vs 32%

production AI adoption, tech
versus education

51%

of the Fortune 500 run an AI
agent in production

2.5

vendors on the average shortlist

Chapter 10. The segmentation framework

Every chapter in this part uses the same three-segment lens, because AI has changed each segment differently (Exhibit 6).

EXHIBIT 6

The three-segment lens

How AI changed each buying environment

SEGMENT	SIZE	COMMITTEE AND CYCLE	WHAT CHANGED
Enterprise	5,000+ employees	13+ internal stakeholders, formal procurement, security review, 6 to 18 months	Every stakeholder now runs their own AI research, and legal reviews AI claims
Mid-market	200 to 5,000	Smaller committees, 1 to 6 months	The most AI-native buyer demographic, with budget and no enterprise approval layers
SMB	Under 200	Owner decision, days to weeks	Self-serve motions and extreme sensitivity to complexity and price

Adoption by segment. AI agent deployment illustrates the gradient (Exhibit 7). Fortune 500 companies show about 51% production adoption of at least one AI agent, with 88% piloting and 3.4 distinct agents on average. Mid-market runs around 34% production adoption, SMB around 22%, and small businesses under 200 employees around 14%. Broader AI adoption follows the same shape, roughly 72 to 78% of large enterprises with AI workloads in production versus 38 to 42% of SMBs, but the growth rates invert: SMB and mid-market adoption is growing faster in percentage terms, and the large-firm gap has been narrowing since mid-2025.

EXHIBIT 7

AI agent adoption by company size

Share running at least one AI agent in production, 2026




The Fortune 500 figure excludes a further 88% running pilots. Source: BCG, Forrester, S&P Global, 2026.

Adoption by industry. Production AI deployment among large enterprises runs from technology and software at about 94% down to education at 32% (Exhibit 8). This gradient predicts both how AI-literate your buyer is and how AI-mediated their research will be. A PMM selling into tech companies faces near-universal AI-first buyers. One selling into the public sector faces committees where AI research is still minority practice.

EXHIBIT 8

Production AI adoption by industry

Share of large enterprises with AI in production, 2026



Source: aggregated enterprise adoption data (Gartner, S&P Global, McKinsey, BCG, Forrester), 2026.

For each vertical below: what changed in the buyer, what changed in the story, the segment plays, and the trap. Playbook guidance synthesizes practitioner patterns and should be pressure-tested against your own win/loss data.

Chapter 11. Financial services and fintech

No vertical combines higher AI adoption with heavier constraint. Financial services sits near the top of enterprise AI production deployment at about 87%, banking and insurance lead agent adoption at about 47%, and marketing in the sector operates inside the most layered compliance surface in B2B: UDAAP at the federal level, FINRA and SEC oversight, and a model-risk framework (SR 11-7 plus the OCC's 2024 generative-AI update) that explicitly treats LLM-based marketing tools as models requiring inventory, validation, and monitoring.

"AI" as a headline claim is simultaneously mandatory and dangerous here. Mandatory because institutions are actively budgeting for it. Dangerous because every AI claim in a regulated communication is an enforceable representation. The winning structure: lead with the governed outcome (fraud loss reduction, disclosure-cycle compression, account-acquisition cost), make the model-risk story a feature (documentation, validation artifacts, audit trails as selling points), and quantify with defensible numbers. Documented compliant-AI programs report account-acquisition cost down about 21% and disclosure cycles compressed about 67% when compliance is designed in from the start.

Enterprise: your AI product will be examined like a credit model. Produce a validation-ready evidence package (model documentation, monitoring approach, bias controls, explainability posture) as core sales collateral. Sell to the second committee: the economic buyer says yes, but model risk, compliance, and vendor risk say when. Anchor proof in peer institutions, because this is the most reference-driven buying culture in B2B.

The trap. Overclaiming autonomy. "Fully autonomous" triggers model-risk escalation. "Human-in-the-loop with audit trail" closes.

Mid-market and fintech: speed with a paper trail. Automated marketing-compliance review is becoming standard infrastructure, and the PMM's job is a claims taxonomy clean enough to automate against. Compete on transparency, publishable pricing, and honest comparison pages, because fintech buyers are heavy AI researchers and thin-content vendors disappear from the answers.

The trap. Affiliate content. Legal responsibility for overpromising affiliates lands on the fintech.

SMB-facing financial products: one-sentence value, visible pricing, instant time-to-value, and social proof from businesses that look exactly like the buyer. AI features should be embedded and invisible ("we chase your invoices") rather than itemized ("AI-powered AR optimization").

The trap. Importing enterprise language downmarket. Compliance vocabulary that reassures a bank alienates a bakery.

Chapter 12. Healthcare and life sciences

Healthcare is the inverse of finance: enormous AI enthusiasm (nearly 90% of healthcare leaders call AI critical) against the lowest tolerance for data mishandling anywhere, producing production adoption around 61% that lags the leaders. For marketers, the defining fact is that the standard B2B toolkit is partially illegal here. OCR enforcement made marketing infrastructure itself a liability: tracking-pixel violations produced settlements, 2024 saw \$9.9M in HIPAA penalties across 22 actions, and standard retargeting, lookalike audiences from patient lists, and pixel-based conversion optimization are effectively off the table for covered entities.

There is no HIPAA-certified AI. Compliance is an operational state, not a product attribute, and sophisticated buyers know it. Credible vendors sell the deployment posture: BAA availability, data-minimization architecture, operation-level audit logging, and honest FDA classification. Tools marketed as "decision support" that function autonomously create FDA exposure for the vendor and liability for the customer.

Enterprise health systems and payers: lead with the governance architecture, because clinical value claims have converged. Peer-reviewed or health-system-validated outcomes beat marketing benchmarks. Champion the clinician: burnout reduction is the emotionally resonant, board-legible outcome of 2026.

The trap. Demo-ware claims about clinical performance, which carry False Claims Act exposure when AI influences Medicare or Medicaid billing.

Mid-market providers and digital health: the fastest adopters and the most frequent compliance casualties. Zero-party data is the strategy: assessments, preference centers, and wellness tools that earn voluntary disclosure outside HIPAA's scope. Content and AEO carry unusual weight because paid channels are constrained, and health queries are among the heaviest AI-assistant use cases.

The trap. Assuming familiar tools are compliant. GA4 by default and Microsoft Clarity are not HIPAA-eligible. A marketing-stack audit is a launch prerequisite.

Small practices: the buyer is a clinician-owner whose top pain is lead follow-up under HIPAA constraints. Sell packaged compliance ("marketing that cannot get you fined") plus visible patient-volume outcomes. Done-for-you positioning wins.

The trap. Feature-selling AI. This buyer buys reduced admin burden and covered liability, not model quality.

Chapter 13. Cybersecurity

Security is the most saturated messaging environment in B2B tech: thousands of vendors, a buyer who is professionally paranoid, and a category where every vendor claims AI on both sides of the equation. Security buyers were early, heavy adopters of AI research tools and are among the most skeptical readers of AI claims. They verify against documentation, community sentiment, and independent testing. The dark funnel is darkest here: practitioner communities shape shortlists no attribution system sees.

Three narrative shifts define 2026 security marketing. From detection to outcomes, quantified: time-to-detect, time-to-contain, analyst-hours saved, false-positive reduction. The agentic SOC story, the category's hottest claim and its most scrutinized: buyers demand to know exactly what the agent does without a human and what happens when it is wrong. And securing AI itself became a category, where early claim-staking with real capability matters and hollow repositioning gets punished by the practitioner community fast.

Enterprise: win the analyst and evaluation layer. MITRE results, analyst placement, and named-logo proof feed both the committee and the AI engines it queries. Build the CISO-peer motion, because peer trust dominates here more than anywhere. The deck the CISO shows the audit committee is a PMM deliverable.

The trap. Fear-based marketing, which reliably backfires with practitioners.

Mid-market: small teams drowning in alerts, the most receptive audience for genuine automation. Sell consolidation and coverage per analyst. Transparent pricing and self-serve evaluation out-convert gated demos. Honest technical content is the highest-ROI channel and doubles as AI-visibility infrastructure, since practitioner write-ups are exactly what the engines cite.

The trap. A "contact sales" wall on pricing sends this buyer to the vendor that publishes it.

SMB: SMBs buy security as an outcome bundled into something else, through MSPs or embedded in platforms they already use. The PMM motion is channel enablement: make the MSP the hero. Direct messaging that works is insurance-adjacent and compliance-adjacent.

The trap. Selling to the SMB as if it has a security team. It has an IT person, or nobody.

Chapter 14. Developer tools and infrastructure

Devtools is ground zero for AI disruption of both the product and the marketing. Developers were always allergic to marketing. Now their first evaluation act is often asking an AI assistant, and their second is asking it to scaffold a proof of concept. This produces the most extreme version of the machine-readability imperative: documentation quality is now discoverability. A tool whose docs

enable an AI to generate working integration code gets recommended and adopted in the same session. And a new layer appeared: AI coding agents themselves influence tool selection by what they default to.

The category's existential question, "does AI make this tool obsolete?", must be answered head-on. Winners repositioned around the AI-augmented workflow, the tool as the thing that makes AI-generated code safe, fast, and deployable. If a developer cannot do something real in minutes, the evaluation is over.

Enterprise: a two-track motion, developer love plus a platform-engineering narrative: make AI-assisted development safe at scale, with policy, provenance, and audit for AI-generated code.

The trap. Leading with the CIO story. Without practitioner pull, top-down devtool sales still fail, and AI has made developer veto power stronger.

Mid-market: product-led growth with human assist. Transparent pricing, usage-based entry, and published benchmarks and honest comparisons, because developers cross-check them with AI instantly.

The trap. Gating docs or benchmarks. Anything gated is invisible to both the developer and the model advising them.

SMB and indie: free tier as marketing, community as channel, template galleries as activation. Indie developers are disproportionately influential in AI-visibility terms: their posts, stars, and forum answers are the citation layer.

The trap. Monetizing too early. The segment's revenue is small but its evangelism sets the models' narrative for everyone upstream.

Chapter 15. Horizontal B2B SaaS

Horizontal SaaS faces the purest version of every dynamic in this report: crowded categories, converging AI features, compressed shortlists, and buyers doing AI-first research. This is where the 2.5-vendor shortlist and the 33% who bought a vendor they had never heard of bite hardest. Review platforms matter twice now, as human validation and as primary AI-engine source material. And the buying committee here is the most AI-native demographic in B2B.

Feature-based differentiation is dying fastest here because AI features replicate in weeks. Durable differentiation moved to data advantage (what your product learns that competitors cannot see), workflow depth, ecosystem position, and category narrative. The agent question dominates roadmap messaging: every horizontal category is being asked when the software becomes an agent that does the work, and PMMs need a credible, staged answer.

Enterprise: the platform consolidation narrative with a migration story. Governance and admin controls for AI features are the enterprise differentiators.

The trap. Roadmap-selling AI agents that do not exist yet. Buyers burned by 2024 and 2025 promises now demand live demos in their own data environment.

Mid-market: the AEO battleground segment. Systematic comparison-page coverage, review-velocity programs, transparent pricing, and content targeted at the exact prompts buyers use.

The trap. Ignoring incumbent-suite gravity. Microsoft, Google, and Salesforce bundling AI features "for free" is the real competitor in most mid-market deals, and the counter-story must be explicit.

SMB: product-led everything, AI embedded not itemized, pricing simple enough to screenshot. 62% of SMB leaders say they will not stay competitive within three years without AI, but 72.5% use a single AI service, so position as the one tool that brings AI into their workflow.

The trap. Trial complexity. Every configuration step before value costs a measurable share of activations.

Chapters 16 to 19. The remaining verticals, in brief

Manufacturing and industrial. Committees pair 30-year plant veterans with digitally native engineers who research with AI, so content must serve both. The narrative is concrete economics under labor scarcity, not "Industry 4.0." Proof is physical: pilots on real lines, reference plants, measurable downtime and scrap outcomes. The trap: cloud-first messaging into OT environments where air-gap requirements are disqualifiers.

Retail and e-commerce. Machine customers arrived first here: roughly a quarter of AI users already use AI shopping assistants, and structured product data is what shopping agents parse. Brands now need to predispose agents as well as people, while authenticity became the human-side differentiator. The trap: personalization claims without a privacy story.

Telecom, energy, public sector, education. Long procurement, political oversight, and the lowest AI-research buyer behavior, so the human committee still rules and channel allocation changes accordingly. Public-sector work is dominated by the compliance story, the incumbent-ecosystem position, and the political-risk narrative. In education, peer districts are the only proof that moves decisions. Energy's 2026 irony: AI datacenter load growth made grid modernization urgent and opened the budgets.

Professional services. AI attacks the billable hour itself, so every purchase is simultaneously an economic decision and a professional-identity negotiation. Peer-firm proof plus workflow-specific accuracy evidence converts. Partners hear "fewer billable hours" as a threat, so the winning frame is capacity for higher-value work under fee pressure. For agencies marketing themselves, the differentiator is a demonstrated AI-native operating model: showing the client how AI is used, governed, and priced into the engagement rather than hiding it.

Insurance. Fairness documentation is now a procurement artifact, and brokers are a mid-market buying wave of their own. Insurtechs should lead with loss-ratio proof, not growth metrics. The trap: claims-handling AI positioned as automation of denials. The winning frame is faster payment of legitimate claims.

Logistics and supply chain. Ops buyers benchmark obsessively and buy against hard baselines. Geopolitical volatility made scenario-planning a first-order purchase criterion. The trap: ROI models built on the customer's data being clean. Discovery that prices in data remediation converts better than optimistic decks that die in implementation.

Real estate, travel, media, HR tech. Common threads: platform-embedded adoption at the small end, integration gravity at the large end, and regulation concentrated where AI touches screening, pricing, or employment decisions. In HR tech, bias-audit documentation moved from differentiator to entry ticket. In media, training-data opacity ends deals. In travel, the strategic conversation is agentic distribution: how the brand shows up when an AI books on a traveler's behalf.

PART IV

The operating model

Team design, governance, and a twelve-month sequence for getting from a 2024 operation to a 2026 one. Chapters 20 to 23.

22% vs 3%

AI marketing budget on generation versus governance

2 to 3

people on an AI-fluent team shipping what six to eight did

1 day

a week protected for judgment work

Chapter 20. Team design, skills, and tooling

Three organizational patterns dominate 2026.

The specialist model. Instead of generalist PMMs each owning a product end to end, teams organize by craft: a positioning specialist, a competitive lead, a pricing lead, an AI-visibility owner, applying deep skill across the portfolio while AI handles the production layer that used to justify generalist coverage.

The PMM-as-hub model. In leaner organizations, PMM absorbs adjacent scope: lifecycle work, website ownership, enablement systems. PMA's data shows sales enablement and website management surging as PMM responsibilities alongside the foundational four.

The fractional and contract layer. A growing share of PMM work runs through specialists hired for defined projects: a pricing overhaul, a category-narrative sprint, a launch. For practitioners this rewards being known for something specific. For leaders it means designing work into contractable units with clean context handoffs.

Team-size math changed. AI-fluent teams of two or three now produce what teams of six to eight did in 2023, but the constraint moved to judgment bandwidth: the number of high-stakes decisions the team can make well per quarter. Headcount cases built on output volume no longer clear. Cases built on decision coverage do.

The 2026 skill stack, in rough priority order: AI fluency as craft (prompt and context engineering, building small automations, knowing when model output is wrong), commercial judgment (pricing, packaging, segmentation, argued with a CFO), evidence discipline (provenance, substantiation, measurement honesty), narrative craft and taste, and systems building: designing the pipelines rather than performing the tasks.

The reference tooling stack changes names quarterly, but the layers do not: a frontier LLM workspace with the team's proprietary context loaded, a research-synthesis pipeline over calls and tickets, competitive monitoring with diff-detection before AI summarization, an AI-visibility tracker across the major engines, a claims governance layer, and workflow automation connecting them. The strategic asset is not the tools. It is the proprietary context, the positioning canon, win/loss corpus, claims file, and ICP definitions that make your instance of the tools better than a competitor's identical instance.


Chapter 21. Governance, risk, and trust

AI-scaled output created a new failure surface, and the budget data says governance is the under-invested layer: organizations spend about 22% of AI marketing budgets on content generation against about 3% on governance, a ratio analysts flag as technical debt in the making (Exhibit 9).

EXHIBIT 9

AI marketing budget allocation

Share of AI marketing spend, 2026



Roughly seven dollars go to generating content for every dollar spent governing it. Source: Improvado and Gartner CMO Spend analysis, 2026.

Claims governance. A canonical, versioned claims file. Substantiation on record for every quantitative claim. Explicit rules for AI-generated customer-facing text, and regulated-industry review lanes. If a regulator asks how a piece of content was approved, "someone reviewed it via email" is not a defensible answer. Timestamped, versioned records are.

Disclosure posture. Norms are hardening around disclosing AI use in content, and, after ads entered AI assistants, around organic-versus-sponsored clarity. Being ahead of disclosure requirements is cheap. Being caught behind them is a trust event.

Data governance. What customer and prospect data feeds which models under which terms is increasingly a sales objection PMM must be able to answer, not just a legal matter.

Brand safety from your own AI. Hallucinated claims in customer-facing AI output are a documented issue at scale, reported by 54% of deploying enterprises. Any AI touchpoint speaking about the product needs the same claims file and review standard as a press release.

The trust thesis underneath all of it: in a market drowning in synthetic content, verified authenticity is appreciating. Real customers, real numbers, real provenance. The 20% of buyers whom AI made less confident are recovered by exactly one thing: credible, checkable evidence.

Chapter 22. A twelve-month roadmap

A sequenced program for a PMM leader starting from a conventional 2024-style operation (Exhibit 10).

The twelve-month roadmap

One layer per quarter, sequenced so each builds on the last

QUARTER	FOCUS	THE MOVES
Q1	Baseline and visibility	Run 30 to 50 buyer-realistic prompts across ChatGPT, Claude, Gemini, Perplexity, and Copilot; log recommendations, framing, and citations; set a share-of-answer baseline against competitors. Build the canonical claims file and kill contradictions across site, docs, and review profiles. Add a free-text "how did you hear about us?" to every conversion point and start mining it. Stand up team AI fluency: a shared workspace, a context library, one automated workflow per PMM.
Q2	The machine-readable foundation	Restructure core pages for extractability. Publish honest competitor comparisons and buyer-prompt-targeted content. Launch the review-velocity and earned-mention program, which is the citation supply chain. Automate competitive monitoring with weekly human-reviewed digests, and retire static battlecards for a living enablement layer.
Q3	The commercial layer	Review pricing and packaging against AI-era value metrics, run willingness-to-pay research, and write the migration story. Rebuild measurement around the triangulation stack and present the new scorecard to leadership. Pick the two industry playbooks closest to your ICP and operationalize them.
Q4	Scale and governance	Formalize claims review lanes, disclosure standards, and AI-output QA. Redesign the launch model around tiers plus always-on machine-visible release communication. Ship the annual narrative: category point of view, analyst cycle, and the board-ready story of what changed.

Throughout: protect one day a week of the team's time for judgment work against the gravitational pull of cheap output. The teams that let AI-enabled volume fill the freed hours end the year busier and less valuable.

Chapter 23. Outlook: 2027 and beyond

Agent-intermediated buying grows but does not complete. Gartner's projection of 90% agent-intermediated B2B buying by 2028 will likely be directionally right for research and shortlisting while human validation persists for high-stakes decisions. The machine-readability and citation work compounds. Start now.

AI answers become a paid medium. Ads inside assistants will create a paid AI-visibility channel and a trust reckoning over organic-versus-sponsored answers. Expect an AEO-plus-paid-placement budget line by 2027, and defend organic credibility jealously. It is the asset the paid layer cannot buy.

Consolidation pressure intensifies. Cheap AI features accelerate platform bundling. Horizontal point solutions face the squeeze first. Category narrative and data and workflow moats become survival equipment, not brand polish.

Governance becomes differentiation. As the EU AI Act phases in and US state law accumulates, the vendors who productized trust convert compliance cost into positioning. Regulated-industry patterns will diffuse into general B2B.

The profession bifurcates. Production-oriented roles compress. Judgment-oriented roles grow in seniority and reporting line. PMA's data already shows PMMs reporting higher in the organization than ever.

AI makes great marketers better and exposes the ones going through the motions. The fundamentals did not change: know the customer, tell the truth compellingly, price the value, arm the field. What changed is that the market, human and machine alike, now checks.

Sources

Primary research and reporting drawn on in this report, accessed July 2026. Where multiple independent sources converge, confidence is higher than for any single estimate.

Forrester, 2026 Buyers' Journey Survey (n of about 18,000), via MR Research, March 2026.

G2, "The Answer Economy: How AI Search Is Rewiring B2B Software Buying," March 2026 survey (n of 1,076), via Demand Gen Report, May 2026.

Gartner: "Future of Marketing: 5 Trends and Predictions for 2026"; 2026 Priorities for Product Marketing Leaders.

Product Marketing Alliance: "What the PMM Role Looks Like in 2026 and Beyond"; State of Product Marketing Report 2026.

Loganix and Averi multi-source analysis (680M citations), PR Newswire, April 2026; Exposure Ninja and SE Ranking conversion data; Yext measurement survey.

Apollo shortlist data and 6sense day-one-shortlist research, via Omnibound and Geisheker Group aggregations, 2026.

Marketing Graham, "How B2B Tech Buyers Research and Buy in 2026" (n of 792), June 2026.

Aggarwal et al., GEO research, SIGKDD 2024; Ahrefs ChatGPT citation analysis, 2025.

Leadscale generational analysis, 2026; Spotlight AR and Profound prompt-volume estimates.

Improvado, "7 AI Marketing Trends for 2026" (budget and governance-gap analysis), May 2026.

Deloitte Digital, Marketing Trends of 2026; Kantar Marketing Trends 2026; Adweek AI marketing analysis, February 2026.

Digital Applied: "Agentic AI for Fintech and Banking Marketing, 2026" (UDAAP, FINRA, SR 11-7 framework); "AI Agent Adoption 2026: 120+ Enterprise Data Points" (Gartner, S&P Global, McKinsey, BCG, Forrester agent data).

Norton Rose Fulbright, Kiteworks, and Anzolo Medical on HIPAA and AI compliance, 2026 (OCR enforcement data).

Presenc AI enterprise adoption statistics, Q1 2026; Goldman Sachs 10KSB, JP Morgan, and Census BTOS small-business AI data, 2026.

BCG and Forrester agent-payback data, 2026; McKinsey State of AI, 2025; Salesforce agentic AI research.

Appendix: key statistics

A consolidated reference of the data points used in this report. Treat single-source figures as directional. Ranges reflect genuine disagreement between studies.

Buyer behavior

METRIC	VALUE	SOURCE
Used AI during their most recent purchase	94%	Forrester, Jan 2026
Buyer AI usage across comparable studies	68 to 89%	Multiple, definition-dependent
Compare vendors inside AI tools	55%	Forrester
Research products inside AI tools	54%	Forrester
Build business cases before vendor contact	47%	Forrester
AI answer engines as a vendor research source	Ranked #1	Forrester
Start software research in an AI chatbot	51%, up from 29%	G2, Mar 2026
Rely on AI chatbots for software research	71%	G2
Chose a different vendor because of AI guidance	69%	G2
Bought from a vendor AI surfaced first	33%	G2
Think more highly of a vendor an AI recommends	85%	G2
Feel more confident in AI-assisted choices	83%	G2
Average vendor shortlist	2.5, down from 3.2	Apollo
Win rate of the day-one shortlist favorite	95%	6sense, 2025
Average buying committee	13 internal, 9 external	Forrester
Prefer sales contact only late in the journey	2 in 3	Forrester
Use AI for supplier research, ages 25 to 34	85%	Leadscale
Use AI for supplier research, ages 55 to 64	23%	Leadscale
Estimated B2B research prompts per day	80 to 100M	Spotlight AR, Profound
B2B buying agent-intermediated by 2028	90%, projected	Gartner

Visibility and AI search

METRIC	VALUE	SOURCE
Website traffic decline as research migrates to AI	10 to 40%	Forrester, Feb 2026
Conversion rate, AI-referred traffic	14.2%	Exposure Ninja, SE Ranking
Conversion rate, Google organic	2.8%	Exposure Ninja, SE Ranking
Extra time on site, AI-referred visitors	68%	Exposure Ninja, SE Ranking

METRIC	VALUE	SOURCE
ChatGPT top-cited pages from DR80+ domains	65.3%	Ahrefs, 2025
Mentions vs backlinks as AI citation predictors	About 3x	Ahrefs, 2025
Citation lift from adding statistics to content	30 to 40%	Aggarwal et al., SIGKDD 2024
Marketers tracking AI visibility	22%	Yext
Unsure how to measure AI search success	64%	Yext
Expect AI to surpass SEO within three years	72%	Yext
Buyer journey visible to traditional attribution	About 27%	Industry estimate

Adoption

METRIC	VALUE	SOURCE
Fortune 500 with an AI agent in production	51%	BCG, Forrester, S&P Global
Fortune 500 piloting agents	88%	BCG, Forrester, S&P Global
Distinct agents per Fortune 500 deployer	3.4	BCG, Forrester, S&P Global
Mid-market agent adoption	34%	BCG, Forrester, S&P Global
SMB agent adoption, 200 to 999 employees	22%	BCG, Forrester, S&P Global
Agent adoption under 200 employees	14%	BCG, Forrester, S&P Global
Median agent-deployment payback	5.1 months	BCG, Forrester
Enterprises running a service agent in production	62%	BCG, Forrester
Average tier-1 ticket deflection	39%	BCG, Forrester
Inference cost decline over three years	About 90%	Industry
Enterprises reporting hallucinated claims in output	54%	Industry
AI marketing budget to content generation	22%	Improvado, Gartner
AI marketing budget to governance	3%	Improvado, Gartner

Industry-by-industry production adoption appears in Exhibit 8.

TWEGS

This report was researched and written by Twegs and published in July 2026. The web version at twegs.com/product-marketing-in-the-age-of-ai is updated as the data moves.

Questions, corrections, or what it means for your team: sophie@twegs.com

Cite as: Twegs, "Product Marketing in the Age of AI: the 2026 report," July 2026.

© 2026 TWEGS · FIRST EDITION · QUOTE FREELY WITH ATTRIBUTION