

THE 2026 MARKET INFLUENCE REPORT

How AI Is Changing the Way Buyers Discover, Trust and Choose Technology Companies

For CEOs, Founders, CMOs, CROs, Communications Leaders and Investors at B2B Technology Companies

AI is not simply changing where buyers search. It is changing which companies enter the consideration set, how those companies are interpreted, and what evidence buyers trust before they ever engage with sales.

Contents

Executive Summary — The Buyer Journey Is Becoming an Answer Journey

Chapter 1 — The New Consideration Set

Chapter 2 — From Search Results to Market Interpretation

Chapter 3 — The Trust Stack

Chapter 4 — Category Control in the Answer Layer

Chapter 5 — Earned Media's New Economic Value

Chapter 6 — Visible Is Not the Same as Believable

Chapter 7 — The Human Validation Layer

Chapter 8 — The Force4 Market Influence Framework

Chapter 9 — What Leaders Should Do Now

Methodology & Sources

EXECUTIVE SUMMARY

The Buyer Journey Is Becoming an Answer Journey

For the past two decades, digital marketing operated on a clear model: earn a search ranking, drive a buyer to a company-controlled website, and convert attention into pipeline. The website was the stage. The brand controlled the narrative. Traffic was the metric.

That model is breaking down.

Buyers are increasingly asking AI systems to do what they once did manually: identify the relevant vendors, summarize the differences, surface common concerns and help build the internal business case for a decision. Those systems draw on media coverage, analyst commentary, customer reviews, community discussions, executive content, technical documentation and hundreds of other signals to construct a picture of the market.

The result is that technology companies are losing control over the first version of their own story. The shortlist is being shaped before a website is visited. The comparison is being framed before a sales call is booked. The credibility judgment is being made before a demo is requested.

The Evidence Is Clear

94% of business buyers now use AI during the purchasing process (Forrester, 2026; ~18,000 global buyers)

51% of B2B software buyers begin their research inside an AI chatbot rather than a search engine (G2, 2026)

92% say AI has influenced their vendor shortlist (Semrush, 2026)

69% have changed their intended vendor choice because of AI guidance (G2, 2026)

33% purchased from a vendor they had never previously heard of (G2, 2026)

89% of unbranded category queries are fulfilled using third-party sources, not brand-owned content (Bain/ScrunchAI, 2026)

84% of all AI citations come from earned media; just 0.3% from paid or advertorial sources (Muck Rack, 2026)

These numbers describe a structural change, not a trend. The buyer journey has a new first chapter, and most technology companies do not yet know it is being written.

The Challenge Is Larger Than SEO

The conventional response to shifts in discovery has been to optimize content. That instinct is understandable but insufficient. Generative Engine Optimization, or GEO, is a real and important discipline. But it addresses only part of the problem.

AI systems do not merely index pages. They interpret them. When a buyer asks an AI tool to recommend vendors in a category, the system must decide which companies belong in that category, how they compare, what their strengths and limitations are, and whether they are credible enough to recommend. That process depends not on the quality of a company's own content alone, but on the breadth and consistency of the external record that surrounds it.

Core Thesis:

GEO is not primarily a content optimization problem. It is a market influence problem. Technical optimization can make content easier to retrieve. It cannot create authority that does not already exist.

AI discoverability depends on the broader record surrounding a company: how media describe it, how customers review it, how analysts categorize it, how executives explain it, and whether credible third parties consistently reinforce the same market narrative.

The Companies That Will Win

The companies that earn consideration, trust and selection in an AI-first buying environment will not simply be those that produce the most content or achieve the highest rankings. They will be those that build what this report calls Interpretation Infrastructure: an interconnected system of positioning, earned media, analyst validation, customer proof, executive visibility and community presence that consistently tells the same credible story.

That system is not new. Its components are familiar. What has changed is that these functions, once managed as separate disciplines, now operate as a single market influence system in which each element either reinforces or undermines every other.

EXECUTIVE TAKEAWAY

A company can lose a deal before it records a website visit, a form fill or a sales interaction. The companies that build deliberate market influence systems will be easier to discover, easier to trust and more likely to be chosen.

CHAPTER 1

The New Consideration Set

How AI Is Reshaping Vendor Discovery Before Traditional Marketing Begins

The most consequential shift in enterprise buying is not happening on company websites or in sales conversations. It is happening earlier, in the AI-mediated research phase that now precedes most vendor contact.

AI as the New First Stop

Forrester's 2026 State of Business Buying report, based on surveys of nearly 18,000 global business buyers, found that AI usage in purchasing decisions rose from 89% to 94% year over year, and that generative AI and conversational search have become the most meaningful single information source cited by buyers — ahead of vendor websites, sales representatives and product experts.

G2's March 2026 survey of 1,076 B2B software buyers confirms the scale of the shift: 51% now begin the purchasing process inside an AI chatbot rather than a traditional search engine, up from just 29% in April 2025. Of those, 63% specifically named ChatGPT as their primary software research tool.

The implications extend beyond discovery. Semrush's April 2026 B2B buyer survey of 643 U.S. professionals found that 55% of buyers use AI to compare vendors, 54% use it to research specific products, and 47% use it to help build the internal business case before making first contact with a vendor.

The Shortlist Is Forming Before You Know It

6sense's 2025 Buyer Experience Report, based on nearly 4,000 B2B buyers, found that the point at which buyers make first contact with a vendor has moved earlier in the journey: from approximately 69% to 61% of journey completion. Buyers are arriving at vendor conversations roughly three and a half weeks sooner than before.

That is not because the journey is shorter. It is because AI has compressed the initial research phase. Buyers arrive at vendor contact having already formed views about the category, identified competitors and, in many cases, made a provisional shortlist.

Key Finding:

G2 found that 69% of buyers changed their intended vendor choice based on AI guidance, and 33% purchased from a vendor they had previously never heard of. AI is not confirming existing preferences. It is actively reshaping them.

Private AI Is Accelerating the Pattern

Forrester reports that 61% of business buyers now use private AI tools provided by their employers. These internal systems often operate on proprietary data and organizational knowledge, creating another AI-mediated layer between buyers and vendor communications that brand teams cannot directly monitor.

Buying Committees Are Larger, Not Smaller

Despite the rise of AI-assisted research, the human dimension of enterprise buying has not simplified. Forrester reports that the typical buying decision now involves thirteen internal stakeholders and nine external influencers. More than 60% of buyers use trials as a risk-reduction mechanism, rising to 78% for purchases over ten million dollars.

AI has not replaced the committee. It has given each committee member a faster path to an opinion, making it more important than ever that the market's description of a company is consistent, accurate and credible across every platform and channel.

Sales Remains Essential, But Its Role Has Changed

Gartner's 2026 research, drawn from a survey of 645 enterprise buyers, is instructive here. While 67% of buyers prefer a sales-representative-free experience for early research, and 70% prefer a fully digital self-service journey, 69% still turn to sales representatives to validate AI-generated insights.

This is the central tension in modern enterprise buying. Buyers want to form their own view without sales interference, but they are not confident enough in AI to act on its recommendations alone. The sales conversation is becoming a validation session rather than an information session. What AI has decided must be confirmed, clarified or corrected.

EXECUTIVE TAKEAWAY

A company can lose a deal before it records a website visit, a form fill or a sales interaction. The new competitive battleground is the AI-mediated research phase that now precedes all of these.

CHAPTER 2

From Search Results to Market Interpretation

Why the Shift From Rankings to AI Answers Is Structurally Different

The decline of traditional search as the primary discovery mechanism for enterprise buyers represents more than a channel shift. It represents a change in what the discovery system actually does.

Traditional Search Presented Sources. AI Interprets Them.

A traditional search engine returned a list of ranked pages and allowed the buyer to form their own view by clicking through, reading and comparing. The brand retained agency over the narrative at the point of engagement. A well-crafted website, a strong landing page, or a compelling case study could still convert a buyer who arrived from a generic search.

An AI answer engine does something fundamentally different. It synthesizes. Before the buyer sees a single vendor's own words, the AI has already:

- Decided which category the query belongs to
- Identified the relevant competitors
- Summarized the primary differences between them
- Surfaced commonly cited strengths and weaknesses
- Made an implicit judgment about which companies are credible enough to include
- Framed the decision in terms the buyer did not choose

A company may never get the opportunity to correct a misclassification, reframe a perceived weakness, or establish its differentiation if AI has already constructed the comparison before the buyer arrives.

The Scale of the Search Shift

Semrush's December 2025 analysis of more than 10 million keywords found that Google AI Overviews appeared on 6.49% of queries in January 2025, peaked at 24.61% in July, then stabilized at 15.69% in November. More significantly, AI Overviews are expanding into commercial, transactional and navigational queries: commercial AI Overview triggers rose from 8.15% to 18.57%, and transactional from 1.98% to 13.94%.

These are not informational queries. These are buying queries. The categories in which technology companies compete are now heavily affected.

McKinsey's August 2025 research found that half of consumers now intentionally seek AI-powered search engines for purchasing research, and 44% of AI-search users describe it as their primary source for buying decisions.

AI Platforms Are Not Identical

A critical error in early GEO thinking was treating AI discoverability as a single environment. Research from Ahrefs, Semrush and Seer Interactive makes clear that the major platforms behave differently:

- **ChatGPT:** ChatGPT cites an average of 15 sources per response (Semrush AI Visibility Index, 2026) and shows weaker correlation with traditional search authority metrics than Google.
- **Google AI:** Google AI Mode and AI Overviews align more closely with classic brand-authority signals, branded anchor text and established trust metrics (Ahrefs Q1 2026 AI Search Benchmark Report).
- **Perplexity:** Perplexity operates its own index with a strong bias toward high-authority, frequently cited publications.
- **Claude:** Claude's web search capability emphasizes grounded, verifiable sources and traceable citations, reflecting Anthropic's product philosophy.

This platform variance matters because a company's visibility profile may look very different depending on where a buyer is researching. Optimizing for one system may not produce visibility in another.

The Citation Economy

Seer Interactive's September 2025 research analyzed 3,119 search terms across 42 client organizations and found that brands cited in AI Overviews experienced 35% higher organic click-through rates and 91% higher paid click-through rates than uncited brands on comparable queries. Citation is creating a compounding advantage: visibility generates credibility, credibility generates engagement, engagement generates signal.

Meanwhile, McKinsey found that only 16% of brands currently track AI search performance systematically. Most companies are competing in an environment they are not yet measuring.

EXECUTIVE TAKEAWAY

The question is no longer simply, 'Can buyers find us?' It is, 'When AI explains our market, does it represent us accurately and advantageously?'

CHAPTER 3

The Trust Stack

The Hierarchy of Signals That Drives AI Visibility and Buyer Credibility

One of the most important and counterintuitive findings in recent buyer research is the degree to which AI systems rely on external, third-party evidence rather than brand-owned content when forming their responses.

Key Statistic:

Bain's April 2026 analysis of approximately 500 million ScrunchAI citations found that 89% of unbranded category prompts are fulfilled using third-party sources rather than brand-owned pages. Muck Rack's May 2026 study of more than 25 million AI-generated links found that 84% of all AI citations came from earned media, with journalism alone accounting for 27%. Paid and advertorial sources represented just 0.3%.

This is not a content quality argument. It is a structural one. AI systems are designed to synthesize from diverse, independent sources. A brand's own website is treated as a primary source for claims, not as evidence. What makes a claim credible is the degree to which independent parties corroborate it.

This creates a hierarchy of trust signals that operates across four layers.

Layer 1: Foundational Signals

The foundation is not about persuasion. It is about legibility. AI systems need to understand what a company is before they can recommend it. This requires:

- Clear positioning and category language that is consistent across the company's own properties
- Entity consistency: the same name, description, category and use-case framing across the website, LinkedIn, Crunchbase, review platforms, directories and documentation
- Technical accessibility: crawlable, well-structured content that can be retrieved and processed by AI systems
- Extractable content: self-contained explanations, direct answers, factual statements and structured data that can be pulled into AI responses without requiring the full page context

Layer 2: Evidence Signals

The evidence layer is where AI systems look for corroboration. These are the signals that move a company from 'mentioned' to 'credible':

- **Earned media:** Business, technology and trade media coverage. Muck Rack's research shows this is the dominant source of AI citations.
- **Customer reviews:** Structured peer validation on platforms such as G2, Gartner Peer Insights and Capterra. G2 describes software review sites as the 'trust layer' in AI-driven software buying.
- **Analyst validation:** Coverage and categorization by Gartner, Forrester and other analyst firms, which carry significant implicit authority in enterprise technology categories.
- **Original research:** Proprietary data, benchmarks and studies that provide evidence other sources can cite and reference.
- **Technical documentation:** Product documentation, knowledge bases and official reference material, which carry particular weight in Claude and other grounded AI environments.
- **Customer proof:** Case studies, testimonials, customer statistics and co-authored content that provide validated evidence of outcomes.

Layer 3: Authority Signals

Authority signals establish broader market credibility and industry standing:

- **Executive thought leadership:** Ahrefs' Q1 2026 AI Search Benchmark Report, analyzing 75,000 brands, found that YouTube mentions had the strongest single correlation with AI visibility. Branded web mentions correlated at 0.66–0.71. Executive-led content on YouTube, LinkedIn and industry publications contributes directly to this footprint.
- **Industry commentary and speaking:** Conference keynotes, podcast appearances and expert commentary in business, technology and trade media establish the intellectual authority that AI systems associate with category leadership.
- **Community participation:** Reddit, LinkedIn communities, GitHub, developer forums and other community platforms are repeatedly cited as sources across major LLM outputs.
- **Wikipedia and reference sources:** Wikipedia appears in the top cited domains across ChatGPT and other systems. Category and company entries in reference sources carry disproportionate weight.

Layer 4: Validation Signals

The final layer is human. Even in an AI-mediated buying environment, buyers complete their validation through people:

- Sales conversations are used to validate AI-generated conclusions, not to receive information the buyer does not already have (Gartner, 2026)
- Analyst and peer recommendations provide the social proof that reduces decision risk
- Trials and demonstrations confirm claims that AI has surfaced but cannot verify experientially
- Procurement and technical reviewers apply the institutional scrutiny that no AI system currently replaces

The Core Principle:

Owned content tells AI what a company claims. Third-party evidence helps AI decide whether to believe it.

EXECUTIVE TAKEAWAY

Companies with thin third-party footprints are structurally disadvantaged in AI-mediated discovery, regardless of the quality of their own content. The Trust Stack is the strategic framework for addressing that disadvantage.

CHAPTER 4

Category Control in the Answer Layer

Why Category Ownership Is Evolving From a Label to an Answer Set

Category ownership has been a central concept in B2B technology strategy for more than a decade. The conventional wisdom is that companies that define and name their category gain a durable competitive advantage: the category becomes associated with the company, and the company becomes the default recommendation within it.

That logic still holds. But the mechanism has changed fundamentally.

From Owning the Label to Owning the Answer Set

In a traditional search environment, a company could assert category ownership through consistent use of a term, investment in content around that term, and enough organic authority to rank for the associated queries. Buyers would encounter the category through the company's own pages and, over time, develop the association the company intended.

In an AI answer environment, the mechanism is different. When a buyer asks, “what are the leading platforms for [category]?”, the AI must answer that question before the buyer visits any company's website. The AI draws on the full external record: analyst coverage, media descriptions, review-platform categorizations, community discussions, and the language used by customers and third-party commentators.

The company that controls the category in this environment is not the one that uses the term most persistently on its own site. It is the one that the widest range of credible external sources consistently associates with the category definition.

Insight:

A category claim that exists only in company messaging is unlikely to survive AI synthesis. A category claim corroborated by analysts, media, customers, communities and peer discussion is structurally durable.

Machine-Legible Positioning

Search Engine Land's entity optimization research highlights a related challenge: AI systems require clarity not just about what a company claims, but about what it is. This means:

- The company's category must be stated consistently across owned and third-party properties

- The use-case fit must be specific enough to be extracted and summarized accurately
- The competitive differentiation must be expressed in language that is recognizable to AI systems trained on the broader market vocabulary
- The entity record across the web must tell a coherent story

Semrush's buyer research adds a buyer-side confirmation: B2B buyers care more about use-case fit than brand recognition in AI answers. The question AI buyers are implicitly asking is not 'who is this company?' but 'is this company right for my specific problem?' Positioning must answer that question.

The Risk of Invented Categories

The traditional category creation playbook encouraged companies to coin new category names that competitors could not easily claim. In an AI answer environment, this strategy carries a structural risk.

AI systems tend to favor consensus language. A category term used only by a single company, without adoption by analysts, media, customers and community participants, is unlikely to appear in AI-generated answers about the market. Worse, the company may be classified under a different, older category term used by its competitors, effectively handing the category framing to the market rather than claiming it.

This does not mean category creation is ineffective. It means it must be paired with an ecosystem adoption strategy: getting analysts, journalists, customers and community participants to use the same language before it can become machine legible.

Competitor Grouping and Category Misclassification

A related and underappreciated risk is competitor grouping. AI systems must place companies in comparison sets to answer vendor questions. If a company's positioning is unclear, AI will make an inference based on the external record. That inference may group the company with low-quality competitors, exclude it from its natural competitive set, or describe it in terms that misrepresent its actual positioning.

Category misclassification is not a theoretical concern. It is an active risk for any company that has pivoted, expanded its platform, or entered a new market segment without updating the surrounding external record to match.

EXECUTIVE TAKEAWAY

Category ownership in an AI-first world requires consistent external corroboration, not just consistent internal messaging. The answer layer is where category battles are now being won and lost.

CHAPTER 5

Earned Media's New Economic Value

How Independent Coverage Has Become Discoverability Infrastructure

Earned media — journalism, specialist publication coverage, analyst commentary and independent review — has always been a component of B2B communications strategy. It has typically been valued primarily for reach, credibility and brand awareness. In an AI-mediated buying environment, its economic function has changed.

The Evidence

Muck Rack's May 2026 Generative Pulse study is among the most significant pieces of primary research in this area. Analyzing more than 25 million AI-generated links across seventeen industries — spanning responses generated by ChatGPT, Claude and Gemini — the study found:

84% of all AI citations came from earned media

27% of AI citations came specifically from coverage published by professional news organizations and trade publications

0.3% of AI citations came from paid or advertorial sources

This pattern has been consistent across three editions of the Muck Rack study since July 2025. It is not an anomaly. It reflects the structural logic of how AI systems are designed: to synthesize from diverse, independent sources, not to amplify brand-controlled claims.

Bain's April 2026 analysis, drawing on approximately 500 million ScrunchAI citations, reaches the same conclusion from the demand side: 89% of unbranded category prompts are fulfilled by third-party sources. When a buyer asks an AI tool 'what are the best platforms for X?', the answer is almost entirely composed of independent sources.

Not All Coverage Is Equal

The strategic implication is not simply 'get more coverage.' The quality, specificity and authority of coverage matter as much as volume.

The most valuable earned media in an AI-mediated environment is likely to be:

- Relevant to the company's category and use cases, not generic brand awareness coverage

- Specific about products, capabilities and differentiation, not vague claims about leadership or innovation
- Published in sources that AI systems recognize as authoritative in the relevant domain
- Consistent with the company's wider positioning across other external sources
- Recent enough (within 12-18 months) to remain retrievable in systems with recency bias
- Supported by data, customer evidence or third-party validation rather than vendor assertions

Generic brand coverage in outlets with no category authority may produce reach. It is unlikely to produce AI citation share in the categories where buying decisions are made.

The Implication for PR Strategy

Communications leaders who have been measured primarily on reach, impressions and media placements are now operating in an environment where the strategic outputs of their work include AI citation share, narrative accuracy, source authority and category consistency.

This is not a minor measurement adjustment. It represents a change in the nature of what effective communications achieves.

The Reframing:

PR has been understood as a downstream awareness function — something that builds reputation after the strategic work of positioning and product development is complete. In an AI-first buying environment, it is better understood as upstream discoverability infrastructure: one of the primary inputs into whether and how AI systems represent a company during the buying process.

Ahrefs' Q1 2026 AI Search Benchmark Report, analyzing 75,000 brands, provides a useful complementary data point. Branded web mentions correlated with AI visibility at 0.66–0.71. YouTube mentions showed the strongest individual correlation — approximately 0.737 — across the entire study. The number of site pages, by contrast, correlated at only 0.194. Volume of owned content is a weak predictor of AI visibility. Cross-web presence and earned brand mentions are much stronger.

EXECUTIVE TAKEAWAY

PR is moving from a downstream awareness function to an upstream input into AI-driven consideration. The companies that build systematic earned-media programs — with category-relevant, specific, authoritative coverage — will accumulate a structural advantage that content volume alone cannot replicate.

CHAPTER 6

Visible Is Not the Same as Believable

Why Citation Share Alone Is an Incomplete Measure of Market Influence

The early conversation around AI discoverability focused almost entirely on visibility: how to appear in AI answers, how to earn citations, how to increase mention frequency. That is an understandable starting point. But it misses a critical dimension.

A company can be visible without being trusted. It can be mentioned without being recommended. It can be cited without being accurately described. And it can appear in an AI answer in a way that actively damages its competitive position.

The Burson Believability Study

Burson's June 2026 research, The Credibility Paradox, provides the most direct evidence of this gap. Partnering with AI marketing platform Profound, Burson analyzed 85 companies across seven AI answer platforms and generated more than 55,000 believability forecasts. The finding: business decision-makers rated AI-generated answers 10% more believable on average than the general population.

Executives are not skeptical of AI. They are credulous. This creates a significant risk: if an AI system describes a company inaccurately, categorizes it incorrectly, or frames it as a secondary option in its category, that description is likely to be believed by the buyers who encounter it. The window to correct the narrative may not open until late in the buying process, if at all.

Six Ways Visibility Fails to Produce Advantage

- **Visible but miscategorized:** The company appears in AI responses but is grouped with the wrong competitors, causing it to be evaluated against an inappropriate comparison set.
- **Cited but inaccurately described:** The company is mentioned but the AI's description of its capabilities, use cases or market position is factually incorrect or outdated.
- **Recommended as a secondary option:** The company appears in the list but consistently in position three or four, with its advantages framed as caveats rather than primary differentiators.
- **Mentioned without credibility signals:** The company is referenced but without the customer evidence, analyst validation or market context that would make the mention actionable for a buyer.

- **Known but excluded from high-value use cases:** The company appears for some query types but is absent from AI responses that match its actual target customer profile and highest-value use cases.
- **Subject to hallucination or outdated information:** The AI produces factually incorrect claims about the company's product, history, customers, pricing or capabilities.

Measuring Beyond Mentions

The measurement frameworks emerging from more sophisticated practitioners move beyond citation volume to include:

- Citation sentiment: is the mention positive, neutral or negative?
- Narrative accuracy: does the AI's description match the company's actual positioning?
- Category inclusion: does the company appear in responses to the prompts that match its target use cases?
- Competitive framing: how is the company described relative to its competitors in direct comparison queries?
- Recommendation rate: in queries asking for recommendations, how often is the company included?
- Audience-specific believability: does the answer resonate with the specific buyer profile it is most likely to reach?

The Reputation Drift Risk

A final and underappreciated dimension of this challenge is what might be called reputation drift. AI systems are trained on and retrieve from large, distributed corpora. If the external record surrounding a company contains outdated coverage, historical controversies, competitor-generated narratives or simply a mismatch between the company's current positioning and its historical description, AI answers may reflect a version of the company that no longer exists.

The discipline of AI-era communications requires not only generating positive, accurate external content but actively monitoring and, where possible, correcting the external record.

EXECUTIVE TAKEAWAY

The goal is not maximum mention volume. It is accurate, credible and favorably framed market representation. Visibility without believability is noise. The competitive advantage lies in being understood correctly.

CHAPTER 7

The Human Validation Layer

Why AI Changes Human Engagement Rather Than Eliminating It

An important counterweight to the urgency of AI discoverability is the persistence of human judgment in enterprise buying. The evidence does not support a narrative in which AI eliminates sales relationships, analyst influence or peer validation. It supports a more nuanced story in which AI reshapes when and why buyers engage with people — and what they need from those interactions.

The Validation Imperative

Gartner's 2026 survey of 645 enterprise buyers captures the central tension: while 67% prefer a sales-representative-free early research experience, 69% still turn to sales representatives to validate the conclusions that AI has helped them reach.

This is not a contradiction. It reflects a logical division of labor. AI provides speed, breadth and pattern recognition in the discovery and comparison phase. Humans provide the contextual judgment, relational trust and risk-management capacity that AI currently cannot. Buyers are using both — in sequence, not in competition.

The Hidden Buyer Problem

One of the most important and under-discussed dynamics in enterprise buying is the role of hidden decision-makers: the internal stakeholders who influence a purchase without appearing in formal buying-group registrations or sales interactions.

Edelman and LinkedIn's 2025 study of nearly 2,000 global professionals found that more than 55% of hidden decision-makers use thought leadership in their vendor vetting process, and 95% say strong thought leadership makes them more receptive to outreach from that company's sales team.

The significance of this finding in an AI context is that hidden buyers are precisely the audience most likely to conduct self-directed AI research. They are researching independently, forming views without vendor contact, and bringing those views into the formal buying process. Reaching them requires being present and credible in the AI-mediated environments where they are forming their opinions.

What Sales Conversations Are Actually For Now

If buyers arrive at sales conversations having already researched the category, identified the competitive options and formed a provisional shortlist through AI tools, the sales conversation has a fundamentally different purpose.

It is no longer primarily an information-delivery event. It is a validation event. The buyer is not asking 'tell me about your product.' They are asking 'confirm the picture AI gave me, address the concerns it raised, and help me de-risk a decision I am already inclined to make.'

This means that sales effectiveness increasingly depends on:

- Understanding which AI-generated conclusions buyers are arriving with
- Knowing which competitors AI tools are introducing into the comparison
- Identifying and preparing for the misconceptions or outdated information that AI may have surfaced
- Being able to provide the human evidence, customer proof and contextual judgment that validates and extends what AI has already told the buyer

Analysts and Peers in a Validation Role

Analyst influence in enterprise technology buying has not faded. It has taken on a second job. Buyers still depend on analyst reports to de-risk decisions, but Forrester finds they are increasingly turning to peer networks to verify what AI search tells them, treating AI-delivered information as incomplete without a human check.

Not all analyst coverage carries the same weight, though. Placements in Gartner Magic Quadrants, Forrester Waves, and IDC MarketScapes generate secondary citation chains: journalists cover them, buyers reference them, peers share them. That compounding is what AI systems pick up.

A placement with a tier-one firm does double duty. It surfaces as high-authority source material in AI-generated buyer responses and serves as the third-party validation buying committees need to justify decisions internally. One asset, two functions. The ROI on the right analyst relationships is higher than most marketing budgets currently reflect.

EXECUTIVE TAKEAWAY

AI is not replacing the human layer of enterprise buying. It is determining what humans are asked to validate. Companies that understand this redesign their sales, analyst relations and customer success programs around validation rather than information delivery.

CHAPTER 8

The Force4 Market Influence Framework

A Five-Part System for Building AI-Era Market Authority

The research presented in this report converges on a single structural conclusion: positioning, PR, analyst relations, reviews, thought leadership, customer proof, documentation and AI optimization are no longer independent marketing disciplines. In an AI-mediated buying environment, they operate as a single market influence system.

A company that invests heavily in one area while neglecting others will find its investment undermined. Strong earned media without clear positioning produces inconsistent AI representation. Clear positioning without third-party corroboration is invisible to AI systems. Original research without distribution fails to enter the external record. Reviews without brand context lack the category signals AI needs to use them effectively.

The Force4 Market Influence Framework describes the five functions that must work together to build durable AI-era market authority.

1. DEFINE

Establish the company's category, market problem, target audience, use cases and competitive distinction with the precision and consistency required for machine legibility. Every external communication — from the website to the analyst briefing to the LinkedIn post — should draw from the same core definition. The goal is not to constrain communication but to ensure that when AI systems synthesize the external record, the output is coherent, accurate and competitive.

2. PROVE

Build the evidentiary layer that third parties can cite and AI systems can retrieve. This includes original research and proprietary data, customer case studies and outcome evidence, analyst validation and category endorsement, technical documentation that demonstrates product depth, and customer reviews that provide structured, searchable peer proof. Proof is not decoration. It is the substance that makes positioning credible to AI systems and the buyers they serve.

3. DISTRIBUTE

Create presence across the channels and platforms that AI systems use as evidence inputs. This means earned media in category-relevant publications, executive thought leadership on LinkedIn and YouTube, analyst briefings, review platform management, community

participation, conference speaking, podcast appearances and cross-web entity consistency. Distribution is not broadcasting. It is ensuring the external record reflects the company's actual positioning across every platform that matters to AI and to buyers.

4. REINFORCE

Ensure the same market narrative appears consistently across the entire ecosystem. The most common failure mode in AI-era brand management is inconsistency: a company's website describing it differently from its analyst profile, its customer reviews using different language from its media coverage, its executive content introducing new category terms not yet adopted by the wider market. Reinforcement is the discipline of maintaining coherence across a distributed, multi-channel, multi-year communication program.

5. MEASURE

Track the signals that matter in an AI-mediated environment, not merely the ones that were important in a traffic-driven world. This includes citation share across major AI platforms, prompt-level competitive inclusion, AI narrative accuracy, source authority distribution, review sentiment and density, branded search trends and validated pipeline attribution from AI-influenced buyer journeys. Measurement informs where the system needs strengthening and provides the evidence base for continued investment.

Framework Summary:

DEFINE your category clearly. PROVE your claims externally. DISTRIBUTE across the channels AI trusts. REINFORCE a consistent narrative. MEASURE what actually drives AI-era consideration.

The Central Strategic Concept: Interpretation Infrastructure

The Force4 Market Influence Framework rests on a concept we call Interpretation Infrastructure: the network of positioning, earned media, analyst validation, customer proof, executive visibility, community presence and technical documentation that collectively determines how AI systems and human buyers interpret a company.

This concept reflects a fundamental shift in how brand strategy works.

In the old model, companies could define themselves on owned surfaces and invite buyers in. A strong website, a persuasive sales narrative and well-positioned content were sufficient to establish identity and differentiation.

In the new model, the market increasingly defines the brand first — through reviews, coverage, documentation, community traces and analyst interpretation — and AI synthesizes that external record into the buying narrative.



Companies that invest deliberately in Interpretation Infrastructure will be easier to discover, more accurately described and more credibly positioned in AI-mediated buying environments. Companies that leave this infrastructure to chance will find that AI fills the gap with whatever the external record contains — accurate or not.

EXECUTIVE TAKEAWAY

Market influence in an AI-first world is not owned. It is built across a distributed ecosystem of external evidence. The Force4 Market Influence Framework is the architecture for building it deliberately.

CHAPTER 9

What Leaders Should Do Now

Audience-Specific Priorities for 2026 and Beyond

The research in this report points to specific strategic priorities that vary by role and perspective. The following recommendations are grounded in the evidence reviewed across more than fifty primary research sources, analyst reports and platform studies.

CEOs

- Treat AI discoverability as market infrastructure, not a marketing tactic. It belongs in the same strategic conversation as product roadmap, channel strategy and competitive positioning.
- Establish executive ownership of category definition. The language used to describe the company, its market and its differentiation must be decided at the top and maintained with discipline across all functions.
- Understand how AI systems currently describe the company. Commission a prompt audit across ChatGPT, Google AI, Perplexity, Claude and Gemini to establish a baseline for category representation, competitive grouping and narrative accuracy.
- Identify reputation and misclassification risks. Understand where the external record diverges from current positioning and create a remediation plan.
- Align marketing, communications, sales and product around a single market narrative. The fragmentation of these functions is the primary structural risk in an AI-mediated buying environment.

CMOs

- Move beyond traffic-only measurement. Add citation share, AI visibility, prompt coverage, competitive inclusion, review sentiment and narrative accuracy to the marketing scorecard alongside traditional metrics.
- Integrate PR, SEO, thought leadership, analyst relations, customer proof and AI visibility into a single program rather than managing them as separate disciplines. They now operate as one system.
- Audit prompt-level competitive visibility. Know which prompts your company appears in, which it is absent from, and what language AI uses when it does include you.
- Strengthen third-party evidence. Review the balance between owned content and earned media, analyst coverage, reviews and community presence. Most companies have over-invested in owned and under-invested in earned.

- Build content strategy around the questions buyers actually ask AI systems. User-intent research must now include AI query analysis alongside traditional keyword and persona research.

Communications Leaders

- Reframe the communications function as discoverability infrastructure, not just reputation management. The primary output of effective PR in 2026 is not impressions — it is AI citation share in category-relevant buyer contexts.
- Prioritize category-relevant, specific, evidence-based coverage over general brand awareness. The most valuable media placements are those that AI systems will retrieve and cite in buyer-facing responses.
- Measure citations and narrative accuracy, not only reach. Track which publications and sources AI systems actually use when describing the company in category-relevant queries.
- Build executive authority deliberately. Executive thought leadership on LinkedIn, YouTube and in specialist publications contributes directly to the brand's AI visibility footprint.
- Track which sources AI systems use. The publications and platforms that appear most frequently in AI responses about the company's category should receive proportionally greater communications investment.

CROs and Revenue Leaders

- Prepare sales teams for the validation role. The buyer who arrives at a sales conversation in 2026 has almost certainly already conducted AI-assisted research. Sales enablement must address this reality.
- Understand which competitors AI tools introduce. Sales teams need to know what comparison sets AI is generating in their category and how to respond to the framing those comparisons establish.
- Identify the misconceptions buyers may bring from AI into sales conversations. Some AI-generated descriptions of products and companies are inaccurate or outdated. Sales teams need to be prepared to correct them without undermining buyer confidence.
- Build proof assets specifically designed for AI-shaped objections. The questions that arise from AI-generated comparisons will increasingly be predictable and should be addressed proactively in sales materials.

Investors and Boards

- Evaluate whether portfolio companies are visible and accurately represented in AI-mediated buying. AI visibility is an early indicator of market position that may precede financial metrics.
- Assess category clarity and third-party validation depth. Companies with weak Interpretation Infrastructure are at greater risk of being excluded from AI-generated consideration sets in their target categories.
- Treat answer-layer visibility as an emerging market moat. Companies that establish early, accurate, credible representation in AI systems serving their buyer categories will compound that advantage over time.
- Identify companies at risk of being excluded or miscategorized. In rapidly evolving categories, the AI representation of the market may lag the actual competitive landscape, creating both risk for incumbents and opportunity for challengers.

McKinsey found that only 16% of brands currently track AI search performance systematically. The companies that act on the insights in this report in the near term will build advantages that late movers will find difficult to replicate. Interpretation Infrastructure takes time to build. The companies that start building it now will have compounding structural advantages — in AI visibility, buyer credibility and market representation — that their competitors cannot acquire quickly.

EXECUTIVE TAKEAWAY

The companies that build deliberate market influence systems will be easier to discover, easier to trust and more likely to be chosen. The window to build that advantage is open now.

METHODOLOGY & SOURCES

Research Basis and Citation Index

This report draws on primary research, analyst surveys, platform studies, academic reviews and proprietary data published between 2025 and 2026. All statistical claims are attributed to their source. Where findings are synthesized or inferred, this is indicated explicitly. All statistics were independently verified prior to publication. The following sources were reviewed in the preparation of this report.

Analyst and Consulting Research

Gartner. 2026 B2B Buying Behavior Survey (n=645 enterprise buyers, published May 2026). Key findings: 45% of B2B buyers used GenAI in a recent purchase; buyers consulted an average of seven information sources; 67% prefer rep-free buying experiences; 69% validate AI-generated insights with sales reps.

Forrester Research. The State of Business Buying, 2026, based on the Buyers' Journey Survey 2025 (published January 21, 2026; ~18,000 global business buyers). Key findings: AI usage in business buying rose from 89% to 94%; generative AI and conversational search are now the single most meaningful information source for buyers; 55% use AI to compare vendors, 54% to research products, 47% to build internal business cases; 61% use private employer AI tools; typical buying group: 13 internal stakeholders, 9 external influencers; 60%+ use trials; AI search often delivers incomplete or unreliable information.

Bain & Company / ScrunchAI. AI Discovery Analysis (April 2026; ~500 million citations analyzed). Key findings: 89% of unbranded category prompts fulfilled by third-party sources; AI is the 'new front door' to unbranded enterprise discovery.

McKinsey & Company. AI Discovery Survey (August 2025); CMO Survey summary (September 2025). Key findings: 50% of consumers intentionally seek AI-powered search; 44% of AI-search users describe it as their primary buying research source; only 16% of brands systematically track AI search performance.

Edelman / LinkedIn. B2B Thought Leadership Impact Study 2025 (~2,000 global professionals). Key findings: 55%+ of hidden decision-makers use thought leadership in vendor vetting; 95% of hidden buyers say strong thought leadership makes them more receptive to outreach.

Platform and Technology Research

G2. The Answer Economy: How AI Search Is Rewiring B2B Software Buying (March 2026 survey; n=1,076 B2B software buyers and decision-makers, published April 15, 2026). Key findings: 51% of B2B software buyers begin research in an AI chatbot (up from 29% in April 2025); 63% use ChatGPT for software research; 69% changed vendor preference based on AI guidance; 33% purchased from a previously unknown vendor. G2 describes review platforms as the 'trust layer' in AI-driven buying.

Semrush. AI Visibility Index 2026; B2B Buyer AI Survey (n=643 U.S. B2B professionals, March–April 2026, published July 2026); AI Overviews Study (December 2025; 10M+ keywords analyzed). Key findings: 66% of B2B professionals regularly use AI to research vendors; 61% compare vendors via AI; 92% say AI has influenced their shortlist; 45% say significantly. AI Overviews peaked at 24.61% of queries in July 2025; commercial query AI Overview triggers rose from 8.15% to 18.57%; transactional from 1.98% to 13.94%; causal zero-click comparison showed modest decline (33.75% to 31.53%) on AIO queries. ChatGPT cites an average of 15 sources per response.

6sense. 2025 Buyer Experience Report (primary study: ~4,000 B2B buyers; supplementary AI/economic uncertainty survey: 766 respondents; published November 2025). Key findings: 94% of buyers use LLMs in the buying journey; vendor interactions remain stable at ~16 per buyer with the winning vendor; first contact moved from 69% to 61% of journey completion; buyers contact vendors ~3.5 weeks earlier.

Seer Interactive. AIO Impact on Google CTR: September 2025 Update (3,119 search terms across 42 organizations; 25.1M organic impressions and 1.1M paid impressions; June 2024–September 2025). Key findings: organic CTR on AIO queries fell from 1.76% to 0.61%; paid CTR fell from 19.70% to 6.34%; brands cited in AIOs saw 35% higher organic CTR and 91% higher paid CTR than uncited brands on similar queries.

Ahrefs. Q1 2026 AI Search Benchmark Report (published May 27, 2026; 13 studies; data including 75,000 brands, 146M search result pages, 730,000 AI responses). Key findings: YouTube mentions are the strongest correlation with AI visibility (~0.737); branded web mentions correlated at 0.66–0.71; number of site pages correlated at only ~0.194; ChatGPT less correlated with traditional authority signals than Google AI Mode.

HubSpot. Answer Engine Optimization Guide and AEO Research (2025–2026). Key findings: schema value debated; AI search creates higher-intent discovery paths; shift from traffic to visibility metrics.

Microsoft. Work Trend Index 2026 (n=20,000 workers across 10 countries). Key findings: 49% of Copilot chats support cognitive work; 66% say AI allows more high-value work time; organizational factors account for 2x+ the AI impact of individual factors; active agents grew 15x year over year.

Salesforce. State of Sales Report (n=4,050 sales professionals, February 2026). Key findings: AI is the top growth tactic; 94% of leaders with agents say they are critical; top performers 1.7x more likely to use prospecting agents.

OpenAI / Harvard / Duke. ChatGPT Usage Study (published September 2025). Key findings: ChatGPT reached 700M users and 18 billion messages/week by July 2025 (~10% of the world's adult population); non-work usage grew from 53% to 73%.

Communications, PR and Reputation Research

Muck Rack. What Is AI Reading? — Generative Pulse Study, May 2026 edition (25M+ AI-generated links across 17 industries; responses generated by ChatGPT, Claude and Gemini; third edition since July 2025). Key findings: 84% of AI citations from earned media; 27% from journalism; 0.3% from paid/advertorial sources. Pattern consistent across all three study editions.

Burson / Profound. The Credibility Paradox (June 2, 2026; n=85 companies across 10 industries, 7 AI answer platforms, 55,000+ believability forecasts; scoring across three audience segments: general population, opinion elites and business decision-makers). Key findings: business decision-makers rated AI-generated answers 10% more believable on average than the general population; visibility without believability is insufficient; fact-based claims are more believable than subjective leadership claims.

Search Engine Land. GEO and Entity Optimization Research (2025–2026). Key findings: AI requires entity clarity; extractable, self-contained content performs better in AI contexts; earned, shared, owned and paid channels all influence generative visibility; top cited domains include YouTube, Wikipedia, Reddit and LinkedIn.

Academic and Independent Research

Academic GEO Literature Review (2026). Critical survey of GEO research. Key findings: strongest evidence supports citation/use improvement conditional on retrieval, not organic discovery; substantial non-overlap between search rankings and AI citation pools; AI systems remain volatile and partially observable.

Harvard Business Review. AI-mediated market discovery and brand positioning analysis (2025–2026). Key finding: discovery, evaluation and recommendation are moving into AI-mediated environments that brands do not own or fully understand.

A Note on Evidence Standards and Fact-Checking

Every statistic in this report was verified against its primary source prior to publication. Sources are cited with their survey methodology, sample size and publication date wherever available. The report distinguishes between observed facts (findings supported by primary research with identified methodologies), emerging trends (directionally consistent findings lacking full causal verification), and inferences (logical conclusions from the evidence not yet directly confirmed by primary research).

The GEO and AI discoverability field is evolving rapidly. Some findings that appear to be emerging trends at the time of publication may be confirmed or contradicted by subsequent research. Readers are encouraged to treat this report as a strategic framework grounded in the best available evidence, and to update their own assessments as new primary research becomes available.