

MARKETING IN THE AGE OF AI

Become the Answer

A competitor can already be winning the deal,
with no click, no form, and nothing in your funnel to show it.

*One company went from 0% AI visibility to
a top-cited authority in under eight weeks.
Here's what actually moved the answer,
and what didn't.*

*Your competitor doesn't need to outrank you anymore.
It just needs to become the answer.*



Somewhere right now, one of your best prospects is asking an AI a question that matters to your business.

Who are the leading providers in this category?
What is the best approach to solving this problem?
Which vendors should a company like ours consider?
How does one solution compare with another?

That prospect may not visit your website.

They may not click an ad. They certainly haven't filled out a form or talked to sales.

But something consequential may already have happened: AI has helped determine who gets considered.

94%

of B2B buyers now use AI during their purchasing process.

Forrester, State of Business Buying, 2026

69%

changed their planned vendor because of AI chatbot guidance

G2, The Answer Economy, 2026

33%

bought from a vendor they'd never heard of before AI recommended it.

G2, The Answer Economy, 2026

For decades, marketers fought for visibility in search results. Ranking mattered because ranking created an opportunity to be evaluated. Buyers searched, clicked, researched, compared, and gradually assembled their own understanding of a market.

AI changes that sequence. Increasingly, it can do some of that work before a buyer ever reaches a vendor's website. AI can explain a category, synthesize information from multiple sources, compare alternatives, identify differentiators, and suggest which companies deserve further consideration. In our earlier work on AI visibility, we described this as the shift from ranking to representation. A buyer can encounter an interpretation of your company before encountering the company itself, and may arrive with assumptions, preferences, and comparisons already formed.

Marketing has always happened before the hand raise. What's changing is how much discovery, interpretation, and comparison can now happen upstream of it.

That makes AI search more than another search channel. It is becoming a new competitive battleground for B2B authority. And your competitor doesn't necessarily have to outrank you to win it. It has to become the answer.

The first wave was about being interpretable. The next is about becoming unavoidable.

The first wave of AI search presented marketers with a relatively straightforward challenge: can AI systems understand who you are, what you're authoritative about, and why you belong in an answer?

That led to an important set of changes.

Companies began paying more attention to clear positioning, content structure, topical depth, consistent terminology, credible sourcing, third-party corroboration, and the relationships between related content.

Those things still matter. They are increasingly becoming the starting point.

Our earlier research also made a broader observation: your website isn't the only thing shaping how AI systems understand your company. AI systems can synthesize information from a distributed digital ecosystem that includes industry publications, comparison sites, partner content, customer discussions, public documentation, and other third-party sources. Your narrative is distributed whether you intend it to be or not.

That realization has helped create an explosion of AI optimization advice. Marketers are told to publish more authoritative content, add FAQs and schema, build backlinks, appear on Reddit, invest in digital PR, create comparison pages, refresh existing content, and monitor hundreds or thousands of prompts.

Much of that advice may be useful. The problem isn't that marketers are doing foolish things. The problem is that an emerging industry has begun presenting tactics as universal answers to a problem that isn't universal.

AI visibility varies by topic, question, competitor, source, and model. Rankings fluctuate.

Methodologies disagree. Attribution is still evolving. The category itself is moving faster than anyone can confidently measure it.

That leaves marketers in an uncomfortable position.

Visibility improves one week and falls the next.

A competitor suddenly appears. A citation disappears.

A new source begins influencing an answer.

*A dashboard can tell you that something changed.
It can't necessarily tell you why.*

Did you cause it?

And if you did, can you do it again?

That's where we believe the next phase of AI visibility needs to go. The question isn't simply, "What are the latest GEO best practices?" It is, "What creates authority in the specific buyer conversations our business needs to win?"

Why most AI content programs stall

None of the common approaches are foolish. Each solves a real problem. Each also has a specific mechanical weakness that shows up once you look at how AI systems actually synthesize an answer.

Four common approaches, and where each one breaks down

The approach	Why teams reach for it	Where it breaks down
Build from a keyword list	Fast, cheap, fits an existing content calendar	Produces isolated facts, not a mapped domain. When AI has to answer a compound or comparative question, there's nothing connective to draw on.
Publish new articles as they're finished	Matches normal editorial cadence, spreads the workload	For months at a time, AI encounters a pillar without its subpillars, or links that point nowhere. Coherence doesn't exist yet during the window the program is trying to earn it.
Add schema and FAQ markup everywhere	Concrete, checklist-friendly, easy to prove you did it	Schema helps a model find structure. It can't manufacture a distinctive answer that isn't there underneath it.
Make every page a pitch	Every asset should "work harder" for the product	Undercuts the independence and third-party corroboration that make content citable in the first place, and it erodes trust with the human who actually lands on the page.

Underneath all four is the same mistake: treating AI visibility as something you add on top of a normal content operation, rather than something that changes what a coherent content operation has to look like.

A dashboard can tell you which of these you're doing. It can't tell you which one is actually working, or why.

Even a marketer who avoids all four mistakes above can still stall, for a different reason: the market has responded to AI search with a wave of new tools, and most of them solve one piece of the problem while leaving the rest for someone else.

Broadly, three kinds of help exist today. Monitoring platforms track prompts, mentions, and citations, but stop at observation; they don't create the content or execute the strategy that would change what they're measuring.

Content platforms help generate AI-assisted writing at scale, but generally don't measure visibility, handle technical optimization, or connect output to an actual AI search strategy. Traditional strategy and consulting firms can plan thoughtfully and bring human judgment, but tend to move slowly, specialize narrowly in AI search, and struggle to operate at the scale a real content and distribution program requires.

We hear the practical consequence from marketing leaders directly: evaluating this landscape can eat up to six months on its own, because the map of vendors keeps expanding faster than anyone can responsibly test them, and no single one of them sits at the center of the loop.

That fragmentation is itself a cost. Every month spent evaluating point tools is a month not spent building authority in the conversations that matter.

This is the gap the Authority Flywheel is built to close: not another tool that observes, or another tool that generates, but a connected operating model that diagnoses the territory, builds the evidence, distributes it, measures what moved, converts it into pipeline, and feeds the results back into the next cycle. One loop, not a stack of point solutions each waiting on the others.

Authority isn't a score. It's a map.

The desire to reduce AI visibility to a single score is understandable. Executives need something they can track. Vendors need something they can put on a dashboard. But a company-wide visibility score can hide the most important part of the story.

We saw this with a large networking and technology company that wanted to strengthen its position across several important security topics. Before the program began, the company had virtually no AI visibility around Zero Trust and Network Security, while its position in adjacent areas varied substantially.

DemandScience used Content-IQ to identify the gaps and build a structured content cluster consisting of one pillar page, four subpillars, and 11 supporting articles. The content was published as a connected body of evidence rather than a collection of independent posts.

Within eight weeks, the company's average AI citation rate reached 14.5%, with an average citation position of #2, while the program generated 18.3K search impressions.

But the aggregate number isn't the most interesting result.

Movement by topic, pre- and post-program

Topic	Before	After 8 weeks	Status
DDoS	1.7% citation rate	25.7% · #2 globally	Strongest win
SASE	23.7% citation rate	24.5% · #6 globally	Increasing
Zero Trust	0.0% citation rate	9.0% · #20 globally	Building
Network Security	0.8% citation rate	3.8%, not yet ranking top	Building

We shouldn't smooth those differences away. They are the lesson.

There is no single thing called a company's "AI authority." A company can be highly authoritative around one problem, moderately visible around another, and almost absent from a third. Each territory can have a different competitive set, different sources of influence, and a different level of difficulty.

The same applies even more strongly to companies with multiple products or business lines. A company may be an established leader in one category, a challenger in another, and a new entrant in a third. Different products may target different buying groups. Some may compete in established categories while others are trying to create new ones. Several products may even contribute to a single customer problem that AI understands differently from the company's internal product structure.

That means the right unit of analysis isn't necessarily the product, business unit, or website section. It's the buyer conversation.

A multi-product company shouldn't begin by asking, "How visible are we in AI?" It should map the commercial territories in which it needs authority and understand its position in each.

A single score tells you whether visibility moved. An authority map tells you where you're winning, where you're losing, and where changing the answer could matter most to growth.

A second experiment: can you enter a conversation you don't already own?

We decided to test that question on ourselves.

DemandScience had established authority around demand generation and B2B performance marketing. But as our business evolved, we wanted to become associated with a newer category and set of conversations where we had little meaningful AI visibility. That made us a useful Customer Zero.

This wasn't an exercise in protecting authority we already had. We were effectively asking whether a company could deliberately enter new competitive territory and begin establishing authority there.

Rather than start with a list of keywords or prompts, we mapped the network of search behavior around the topic: what people were asking, how those questions connected, which areas represented meaningful opportunities, and how content could be structured around that network.

That approach is central to Content-IQ: understanding search behavior as a connected network rather than a collection of isolated keywords.

We built that map into a coordinated 16-asset cluster (one pillar, five subpillars, ten supporting articles) and published it as a single, simultaneous launch rather than a drip of individual posts.

We then measured what happened across AI engines against the competitors already established in that conversation.

The movement came much faster than we expected.

Within about two weeks of launch, DemandScience went from no meaningful presence in the category to a 10% overall mentions share across tracked AI engines, good enough for third place among the five vendors we tracked, ahead of two established competitors in the category.

Mentions share vs. tracked competitors, by AI engine

(~2 weeks post-launch · source: internal AI search tracking)

Engine	DemandScience	Vendor A	Vendor B	Vendor C	Vendor D
All Engines	10%	39%	40%	5%	2%
ChatGPT	13%	31%	31%	6%	1%
Perplexity	7%	28%	30%	3%	1%
Google AIO	6%	46%	49%	5%	2%
Gemini	7%	55%	48%	1%	3%

Competitors anonymized as Vendor A-D. No competitor is named anywhere in this piece.

The breakdown by engine is where it gets more specific. DemandScience's strongest showing wasn't the average, it was ChatGPT, where mentions share reached 13%, ahead of its own overall number and closer to the two established leaders than the aggregate figure suggests.

AI authority may not be predictable, but it is movable, and it can move faster than most marketing timelines assume.

Two established competitors still lead the category overall. That's not the story. The story is that a company with zero presence in a conversation reached third place, ahead of two other established vendors, in about two weeks.

The point isn't that we now own the category. We don't, and closing the gap with the two largest players will take sustained work, not one launch. Nor does one experiment give us a universal formula for creating AI authority.

We now had two different examples. In one, an established technology company dramatically increased authority in some security conversations while others moved more gradually. In the other, DemandScience entered territory where it had essentially no established AI presence and reached a ranked, meaningful position within weeks.

That changes the problem for marketers. If authority can move this quickly, the next question isn't simply how to monitor it. It's how to deliberately influence it, learn from what happens, and determine where that influence creates business value.

From visibility monitoring to an Authority Flywheel

We believe that requires a different operating model. Rather than treating AI visibility as an optimization project or a new dashboard, companies should approach it as a continuous cycle.





01. Diagnose

Find the battles worth winning

Identify the buyer conversations with enough commercial significance to deserve investment, not every conceivable prompt. For multi-product companies, this includes portfolio prioritization: which categories are strategic bets, where you're already strong, and where the market is open enough to plausibly win.



02. Build

Give the market evidence you belong in the answer

More content isn't the goal. Depth, coherence, credibility, and connection are. A structured cluster beats a pile of unrelated articles, and a multi-product company needs one coherent knowledge architecture, not five independent AI strategies.



03. Amplify

Authority can't live only on your website

Earned media, executive thought leadership, customer advocacy, partners, communities, social, events, syndication, and paid all extend the reach of the evidence you've built. We're not claiming a fixed causal formula here — the honest answer is experimentation: build, distribute, observe, repeat what works.



04. Measure

Find out what moved it, not just whether it moved

Citation rate, share of voice, and average position are observation, not understanding. The valuable question is what you did immediately before the number changed, and whether that improvement persisted.



05. Convert

A citation isn't the business outcome

The question is whether greater authority changes buyer behavior: traffic quality, engagement, and ultimately pipeline. For a multi-product company, connecting authority-by-topic to commercial outcomes tells you where to actually invest.



06. Learn

Build an advantage nobody can sell you in a dashboard

Which interventions were associated with movement? Which gains persisted? Feed those observations back into the next diagnosis. The goal isn't a permanent formula — it's a marketing organization that learns faster than the environment changes.

Diagnose: find the battles worth winning

The starting point isn't monitoring every conceivable prompt. It's identifying the buyer conversations with enough commercial significance to deserve investment.

Which questions shape how buyers understand your category? Which ones influence consideration? Where are competitors becoming the default answer? Where is your company absent, misunderstood, or represented in ways that no longer reflect your strategy?

For a company with multiple product lines, diagnosis also requires portfolio prioritization. Which categories are strategic growth bets? Where are you already strong? Where is the market open enough that authority can plausibly be gained? Which products should reinforce a larger solution narrative rather than compete for independent visibility?

This is why thousands of prompts aren't necessarily an advantage. More measurement can simply produce more noise.

Instead of asking, "How do we rank better in AI?" the question becomes, "Where would changing the answer actually matter to our business?"

The goal should be a prioritized map of commercially meaningful conversations. For each territory, marketers should understand who appears today, how competitors are characterized, what sources appear to support the answers, and what gaps exist between the authority the company has and the authority it wants.

Build: give the market evidence that you belong in the answer

Once the territory is understood, companies need enough evidence to establish authority within it. Content is central to this, but more content isn't necessarily the answer. AI systems already have an enormous amount of material available to them. Another generic article doesn't automatically make a company more authoritative.

The objective is depth, coherence, credibility, and connection.

Our networking-company example illustrates the distinction. Rather than publishing 16 unrelated articles, the program created a structured cluster with a pillar, subpillars, and supporting content designed around the territory the company wanted to strengthen.

For multi-product companies, coherence becomes especially important. Product teams often use different language for related problems. Acquired companies bring legacy terminology and content. Different business units may make overlapping or even contradictory claims. Products that are supposed to work together can appear unrelated when their digital evidence is viewed from outside the company.

AI doesn't necessarily understand the organizational reasons for those inconsistencies. It has to synthesize the evidence it encounters. As our first AI visibility paper argued, AI doesn't resolve inconsistencies; it can blend them, allowing fragmented positioning to become part of how a brand is represented.

A company with five products therefore doesn't necessarily need five independent AI strategies. It needs a coherent knowledge architecture that makes clear what problems the company solves, where each offering fits, how the offerings relate, and what larger areas of expertise the company can credibly own.

What this looks like in practice

Here is how we actually build a cluster, using our ABM program as the working example.

Start with the question network, not a keyword list.

Before writing anything, we map how the questions in a topic actually relate to each other: which ones a buyer asks first, which follow from those answers, and where the connective tissue between them lives. That map, not a spreadsheet of search terms, determines what gets built.

Size the architecture by tier, deliberately.

A pillar page is built to carry the small number of questions that define the category itself. Each subpillar goes one level deeper into a major sub-question. Each supporting article goes deeper still, into a single decision a buyer has to make. The same shape has now shown up twice for us: our security-industry cluster used one pillar, four subpillars, and eleven supporting articles; our own ABM cluster used one pillar, five subpillars, and ten supporting articles. Different topics, same underlying architecture.

Wire the link graph before anything publishes.

Every supporting article links up to its subpillar; every subpillar links up to the pillar and down to its supporting articles. We also build deliberate cross-links between subpillars that a reasoning model, or a reasoning buyer, would naturally connect: a measurement subpillar links to the one on common pitfalls, because anyone asking about ROI is also implicitly asking what makes ROI hard to prove. In our process, broken or orphaned links at launch are treated as a hard failure, not a fast-follow fix.

Launch the architecture at once, not piece by piece.

The full cluster, pillar and all subpillars, goes live on the same date. This is the direct fix for the drip-publishing problem above: there is no multi-month window where AI systems encounter half a structure.

Retire the overlapping legacy content at the same moment.

Where old, thinner posts covered ground the new cluster now covers better, they get redirected into the new pages rather than left standing alongside them. In the ABM program, one new article on account tiering formally replaced three older, thinner posts on the same topic, live at launch, not as later cleanup.

Calibrate how much of each page is "about us."

Educational, strategy-level content stays vendor-neutral, because that's where trust gets built, and where being visibly promotional would cost more credibility than it buys. We lean into our own story only where it's a genuine differentiator, or where candor itself is the differentiator: an entire subpillar in the ABM cluster is built around why ABM programs fail, and another around when ABM is the wrong move for a given company. Telling a buyer when not to buy is a credibility signal an all-promotional competitor can't credibly make.

*The mechanism isn't a secret.
It's discipline: map the questions,
size the architecture on purpose,
wire it before it's live, launch it whole,
retire what it replaces, and tell the truth
about the parts of the story that
aren't flattering.*

Amplify: authority can't live only on your website

Building authority is necessary. Leaving it on your website and waiting for the market to discover it is unlikely to be sufficient.

AI systems don't interpret companies solely through their owned properties. The broader digital ecosystem matters. Third-party publications, comparison sites, partners, customers, communities, public documentation, and other sources can contribute to how a company and category are represented.

If authority is distributed, marketers have to think about how the ideas and evidence they build get distributed through the market. That can include earned media, executive thought leadership, customer advocacy, partners, communities, social, events, content syndication, paid social, display, and other forms of amplification.

We should be clear about what that does not mean.

We aren't claiming that running a LinkedIn campaign directly causes an LLM to cite your company, or that a particular level of paid-media exposure produces a predictable increase in AI share of voice. There isn't credible evidence for a universal formula like that today.

In fact, that uncertainty is part of the opportunity. Nobody can credibly give a CMO a fixed weighting of all the signals that influence AI representation across every topic.

The answer shouldn't be another universal checklist. The answer should be experimentation.

Build credible authority around a commercially important conversation. Put that authority into the market through the channels available to you. Observe what changes. Repeat the interventions that appear to work and stop investing in those that don't.

This also gives DemandScience a broader role than a pure AI visibility or GEO vendor. Content can build the foundation of authority, while multi-channel activation can help put those ideas into the market.

If AI learns about a category from the market, your authority has to exist in the market.

Measure: don't just watch the score. Find out what moved it.

AI visibility platforms are useful. They can tell marketers whether a company appeared, how frequently it appeared, which competitors appeared alongside it, and which sources were cited. But that's observation. It isn't yet understanding.

The more valuable question is what happened after an intervention. Did the targeted topic move? Did adjacent topics move? Did the competitive set change? Did different sources begin appearing? Did the improvement persist? What did the company publish, distribute, earn, or change immediately before the movement?

The technology-company example shows why this matters. A 14.5% average citation rate sounds impressive. But DDoS at 25.7%, Zero Trust at 9.0%, and Network Security at 3.8% provide far more useful information about where authority was actually being created and where additional work remained.

Perfect attribution may not be possible.
Models change. Competitors act independently.
Some signals are difficult to isolate. We shouldn't confuse correlation with proof of causation.

There is a large difference between saying "our visibility changed" and saying "we made a deliberate intervention, established a baseline, observed what happened, and now have evidence about what appears to move it."

Convert: a citation isn't the business outcome

Citation rate increased. Share of voice went up.
Average position improved. The dashboard turned green.
Now what?

The business question is whether greater authority changes buyer behavior.

60.1%

AI-referred
engagement rate
vs. 40.8% organic

1:08

avg. engagement
time vs.
0:31 organic

10.72

events / session
vs.
7.64 organic

4.05%

session key-event
rate vs.
0.82% organic

We shouldn't overstate what those numbers prove. AI traffic remains a developing channel, and one company's traffic patterns don't establish a universal rule. But there is a plausible reason the behavior could be different: a buyer arriving after an AI interaction may already have asked several questions, explored alternatives, compared approaches, and developed a clearer idea of what they need.

Our more recent Customer Zero experience provides another signal. As DemandScience's AI mentions increased around targeted topics, we saw more high-intent website leads coming from AI, and those leads converted at more than three times the rate of organic traffic.

For a multi-product company, connecting those dots becomes even more useful. One category might produce large amounts of AI visibility but little commercial engagement. Another might generate less traffic but much stronger buying behavior.

That's when AI visibility stops being another interesting marketing metric. It starts looking like a potential pipeline lever.


Learn: build an advantage nobody can sell you in a dashboard

Which topics moved? Which didn't? Which interventions were associated with the movement? Which third-party sources became more important? Which gains persisted? Which authority improvements generated engagement?

Those observations should feed the next diagnosis. Over time, the company develops a proprietary body of evidence about what appears to create authority in its markets, against its competitors, with its buyers. That's more valuable than knowing the industry's current list of GEO best practices.

The battleground won't remain static. Models will change. Sources will change. Competitors will respond. Buyer language will evolve. New categories will emerge.

The objective isn't to discover the permanent AI ranking formula. It's to build a marketing organization capable of learning faster than the environment changes.



The invisible buying cycle just became more invisible

None of this changes one of the oldest truths in B2B marketing: much of marketing's job happens before a prospect is ready to talk to sales.

AI doesn't replace that behavior. It changes how efficiently, and how invisibly, some of it can happen.

A buyer can now ask an AI system to explain a market, identify possible approaches, recommend companies to investigate, summarize strengths and weaknesses, and compare alternatives.

For a multi-product company, the implications are even larger. AI can become an interpreter of the portfolio itself.

It may determine which of your products is relevant to the buyer's problem, whether several offerings belong together, how your broader solution compares with a specialist competitor, or whether your company belongs in the category at all.

This is why we don't believe the strategic issue is ultimately GEO, AEO, LLMO, or whatever acronym comes next. The larger issue is that the mechanics of B2B consideration are changing.

What would happen if you became the answer?

Don't start with 10,000 prompts. Start with the questions that could change your business.

Ask who AI recommends today. Look at how those companies are characterized and what evidence appears to support those answers. Determine where your company is already authoritative, where it's absent, and where it's being misunderstood.

Then make a choice. Which of those territories are actually worth fighting for?

Build the evidence required to compete there. Amplify it across the market. Establish a baseline. Make deliberate interventions. Measure what happens.

Learn from the results. And don't stop when the visibility number moves — ask whether buyer behavior moved with it.

Buyers haven't stopped buying. They've changed how they discover, research, compare, and decide which companies deserve further consideration. They're already asking the questions. AI is already answering them. Your competitors are already appearing in those answers.

The pipeline is out there. The question is whether the answers lead buyers to you.

Your competitor doesn't need to outrank you anymore. It just needs to become the answer.

Diagnosis Determines Everything.

Understand which buyer conversations matter, where competitors are winning, and what's actually winnable.

Get Your Market Intelligence



DemandScience is a Precision Performance Marketing company that helps B2B organizations generate predictable pipeline by connecting verified buyer signals, high-impact content, and expert multi-channel execution. Trusted by 900+ leading B2B companies, DemandScience delivers an integrated, managed system that eliminates the operational burden of fragmented marketing stacks and replaces activity metrics with measurable revenue outcomes. Visit demandscience.com.