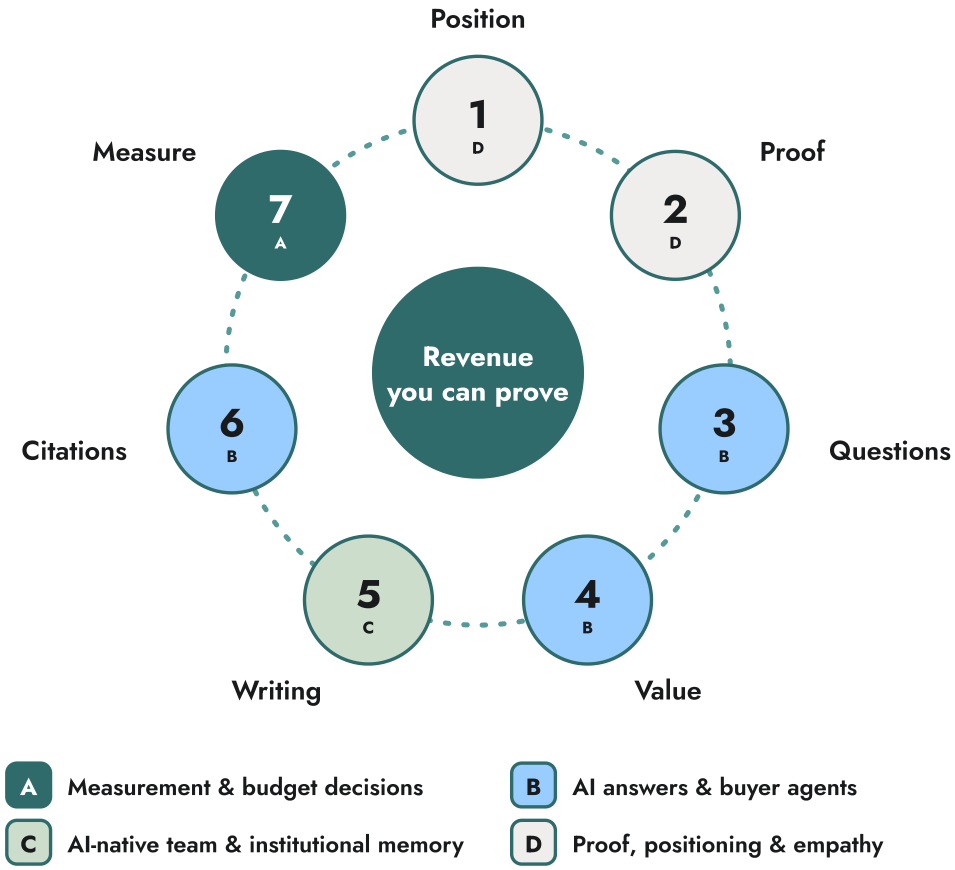
PMB · PAUL BUCKLEY

The Definitive Guide to Marketing in the Age of AI (in Seven Steps)

AI now answers your buyers' questions, writes your competitors' content and decides which brands get named. These seven steps still grow revenue.

pmbdirect.com/the-definitive-guide-to-content-marketing · Rewritten for the AI era
October 26, 2026



THE GUIDE

The Definitive Guide to Marketing in the Age of AI (in Seven Steps)

AI now answers your buyers' questions, writes your competitors' content and decides which brands get named. These seven steps still grow revenue, rebuilt for buyers who ask machines first.

Originally published February 24, 2025 as The Definitive Guide to Content Marketing. Rewritten for the AI era October 26, 2026.

Marketing in the age of AI comes down to seven steps. Own a position others can repeat. Back it with proof a machine can quote. Plan around the questions buyers and their agents ask. Do the buyer's work. Write for three kinds of reader. Earn citations instead of chasing clicks. Measure what grows revenue when the clicks disappear. Content got cheap. A clear position, checkable proof and honest measurement did not.

> **In short.** AI didn't kill marketing fundamentals. It made them the only thing that separates you. When every team can publish ten times more, the brands that win are the ones AI names, buyers remember and CFOs can see paying off.

I wrote the first version of this guide for buyers who searched, clicked and read your site. Many of them now stop at the AI answer. This is the rewrite.

What changed for marketers when AI started answering first?

Three things changed at once. Buyers ask machines first, machines write most of the first drafts, and the click that held measurement together is disappearing.

Buyers ask machines first. Google says AI Mode passed one billion monthly users within a year of launch (Google). In B2B, 45% of the 645 buyers Gartner surveyed used generative AI in a recent purchase. And 69% prefer to check what the AI told them with a sales rep (Gartner).

Machines write the first draft. In the Content Marketing Institute's research for 2026, 87% of B2B marketers using AI for content say productivity improved. Only 39% say content performance did (CMI).

The click is disappearing. Pew Research Center tracked 68,879 real Google searches. When an AI summary appeared, people clicked a regular result on 8% of visits. Without one, they clicked on 15% (Pew). Ahrefs found that an AI Overview goes with a 58% lower click-through rate for the top-ranked page (Ahrefs).

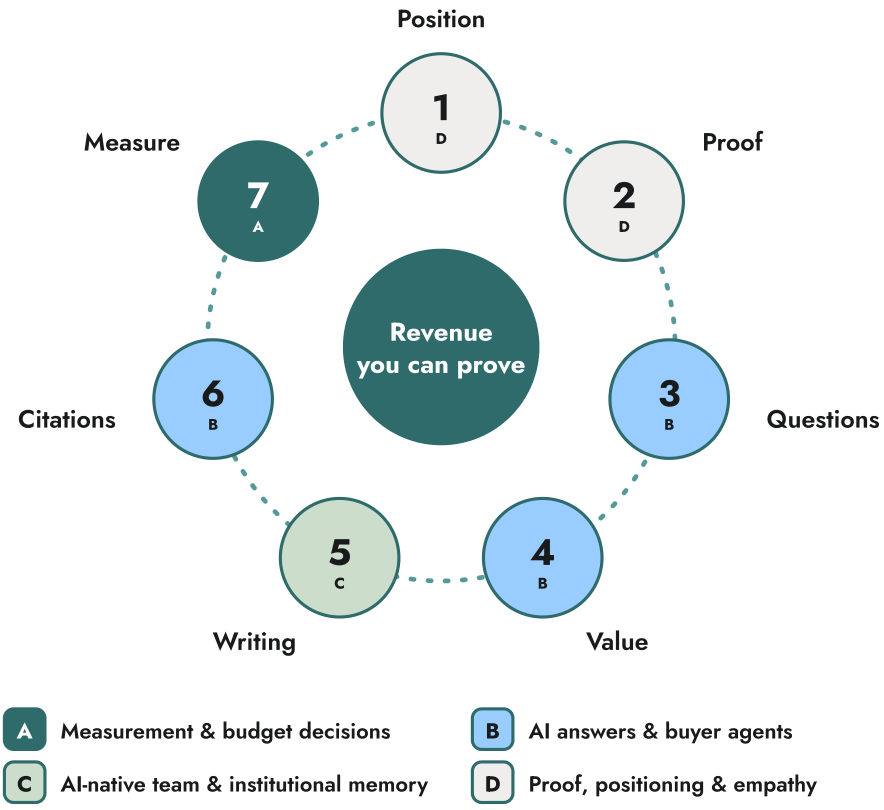
Put those together and you get the trap most teams are in: more content, same results.

What are the seven steps of marketing in the age of AI?

The seven steps are the same spine as the original guide, but each one now answers a different question. The table shows what each step used to ask and what it has to ask now.

The seven steps, then and now

Step	The original question	The question now	Chapter
1. Position	How do we stand apart on our own site?	Will AI and other people repeat our position in one sentence?	Position for the answer
2. Proof	What are our message pillars?	Is every claim specific, sourced, dated and quotable?	The proof ladder
3. Plan	What do we publish each month?	Which buyer and agent questions do we own, one page each?	The question map
4. Value	How do we help buyers on our site?	Does our content do the buyer's work, and their agent's?	Do the buyer's work
5. Write	How do we persuade a human reader?	Does it work for the skimmer, the scanner and the machine?	Write for three readers
6. Distribute	How do we drive traffic?	Where do AI answers get their sources, and are we there?	From traffic to citations
7. Measure	Which headline wins the A/B test?	What grows revenue when there's no click to count?	Measure what grows revenue



The seven steps form one loop. The letter in each circle is the job it serves: A measurement, B AI answers, C the AI-native team, D proof and positioning.

Which step should a CMO start with?

Start with step 7 if you own a budget, and with step 1 if you own a brand. Most CMOs own both, so read chapter 7 first and chapter 1 second.

Here's why measurement goes first. Every other step costs money, and AI makes it easy to spend that money on things that are simple to count. If you can't show what grows revenue, the budget drifts toward whatever reports best. Gartner's 2026 CMO Spend Survey found budgets flat at 7.8% of company revenue, with 56% of CMOs saying they lack the budget for their strategy (Gartner). There's no slack for guessing.

Where to start, by the problem you have

If your problem is…	Start with	Then
“I can’t prove what our spend is growing.”	Chapter 7	Chapter 6
“AI answers name our competitors, not us.”	Chapter 6	Chapter 1
“Our content sounds like everyone else’s.”	Chapter 1	Chapter 2
“Sales says marketing content doesn’t help close deals.”	Chapter 4	Chapter 2
“My team uses AI, but output went up and results didn’t.”	Chapter 5	Chapter 3

How do the seven steps fit together?

The steps form one loop across four jobs. Steps 1 and 2 build what’s worth saying: position and proof. Steps 3, 4 and 6 get you into the answers buyers and their agents read. Step 5 is how an AI-native team produces the work without losing judgment. Step 7 closes the loop, so the next cycle starts smarter than the last.

> **The four jobs behind the seven steps.** 1. **Measurement and budget decisions** (step 7): know what grows revenue, and fund it. 2. **AI answers and buyer agents** (steps 3, 4 and 6): be the source AI names. 3. **The AI-native marketing team and institutional memory** (step 5): staff AI like a team, and write down what you know. 4. **Proof, positioning and empathy** (steps 1 and 2): say something specific, true and human.

Every chapter ends with a **scoreboard**: how you’ll know the step worked. The old guide saved testing for the last step. This one plans the measurement on page one.

What stays the same?

People still decide on feeling and justify the decision with reasons afterward. They still trust specific stories over vague claims, and they still buy from people who understand their problem. Those ideas carried the original guide, and they carry this one. What changed is who reads first. Now it’s often a machine, and the machine passes along only what’s specific enough to quote.

That’s also why empathy matters more now. When every brand can generate fluent copy, the scarce thing is knowing what your buyer fears and hopes for. I wrote about that in The Power of Empathy and Why We Like Stories.

What doesn’t work anymore?

Three habits from the old playbook now cost you money.

1. **Publishing for volume.** More pages split what little authority you have across near-duplicates. AI engines quote one passage per question. Give each question one owner page.
2. **Counting traffic as success.** Traffic is the wrong scoreboard when the answer appears before the click. The win is being named, cited and chosen.
3. **Testing small things while the big decision goes unmeasured.** A headline test is fine. It won't tell you whether to move a million dollars from one channel to another. That takes mix modeling and real experiments.

Who wrote this guide, and why trust it?

I'm Paul Buckley. I've worked in marketing for 35 years, since 1991, and founded PMB in 2005. Every result below predates generative AI. The skills behind them (a clear position, checkable proof and honest measurement) are the ones AI makes scarce.

- **Dell.** At RAPP (1991–2005), the consumer programs I led generated over $1.5 billion in annual revenue for Dell.
- **Harrah's.** In 2012, working in partnership with Polaris Research, we used marketing mix modeling to cut more than $30 million a year from marketing spend at a single Harrah's property, with performance maintained.
- **D Custom.** Where I was president from 2014 to 2024, revenue doubled and client work earned more than 60 national awards.

Disclosure: Paul is EVP, Business Development at Polaris Research, the modeling partner on the Harrah's work.

I also co-founded Berry Buckley with Drate Berry. More about my background is on About Paul and Selected Work.

Where should you go from here?

Pick the chapter that matches your problem, or read them in order. Each one stands on its own. If you'd rather work through the seven steps with a senior partner beside your team, talk to me about advisory support: work with me.

The teams that win this shift won't publish the most. They'll be the ones machines quote and buyers remember.

FAQ

What is marketing in the age of AI?

It's marketing for buyers who ask AI first. The fundamentals (position, proof, value, measurement) still decide who wins, but the first reader is often a machine that quotes only specific, sourced passages, and the click that measurement relied on is disappearing.

Is content marketing dead because of AI?

No. Lazy content is. AI made drafting cheap, so volume stopped being an advantage. Content that states a clear position, carries checkable proof and answers real buyer questions still earns citations, sales use and revenue.

What should a CMO do first about AI search?

Fix measurement first, so you can see what grows revenue when clicks fall. Then check whether AI answers name your brand for your buyers' top questions. Chapter 7 and chapter 6 cover both.

How is this guide different from the original version?

Same seven-step spine, new questions. Each chapter now plans for AI answers and buyer agents, uses current sourced research and ends with a scoreboard that shows whether the step worked.

Is there a PDF of the guide?

Yes. The full guide is a free PDF download with no form to fill in. You can leave your email to get one note when the guide is updated, but you don't have to.

> Read online: https://www.pmbdirect.com/the-definitive-guide-to-content-marketing

STEP 1 OF 7

Position for the Answer: Be the Brand AI Names and Buyers Remember

Brand positioning used to be about standing apart on your own website. Now it has to survive being retold by journalists, customers and the AI models that brief your buyers.

Step 1 of 7 in The Definitive Guide to Marketing in the Age of AI (in Seven Steps). Originally published October 11, 2024. Rewritten for the AI era October 26, 2026.

Brand positioning in the age of AI has two audiences: the buyer, and the machine that briefs the buyer. A position works when it's **aligned** with what motivates buyers, **apart** from what competitors claim and **repeatable** in one sentence by someone who isn't you. If a journalist, a customer and an AI model can't restate it, you don't have a position yet.

> **In short.** Aligned, apart, repeatable. The first two came from the classic perceptual map. The third is new, because AI mostly repeats what other people say about you.

Why does positioning matter more when AI writes everyone's copy?

Because the default output of AI is the average of everything it has read. Ask a model to describe a B2B services firm and you'll get "trusted partner," "innovative solutions" and "customer-centric approach." Your competitors are getting the same sentences.

That isn't a guess. In CMI's research for 2026, 89% of B2B marketers whose organizations use AI apps use content creation tools that generate or optimize marketing copy (CMI). In HubSpot's 2026 survey, 62.7% of marketers said they need more unique, human-centered content to compete with AI content (HubSpot). When everyone's drafts come from the same engine, a sharp position is the only thing a reader can tell apart.

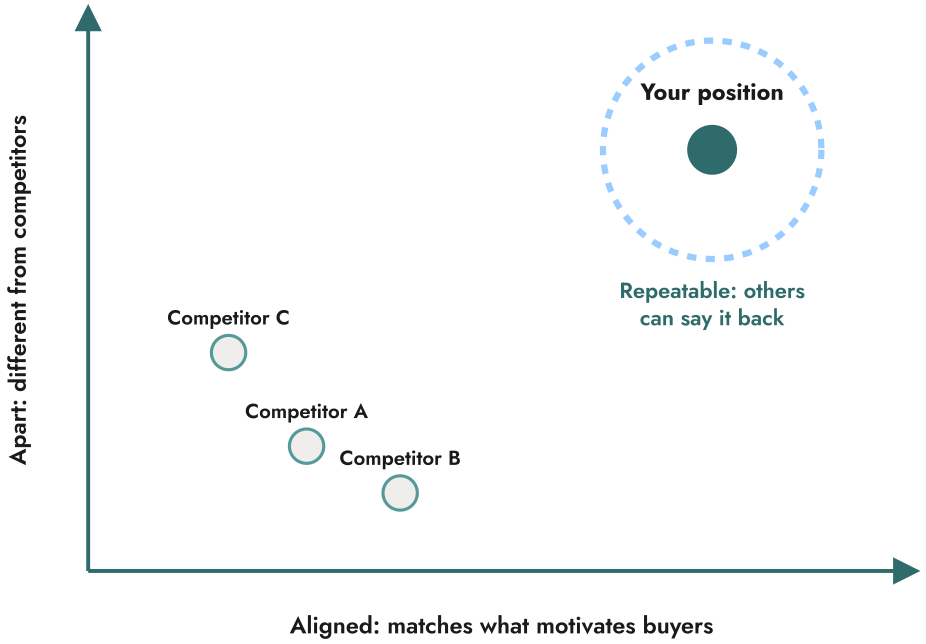
Where does AI get its picture of your brand?

Mostly from other people. Muck Rack studied more than 25 million links cited by ChatGPT, Claude and Gemini in May 2026. Earned media, meaning journalism, research and other people's sites, made up 84% of the citations. Paid content made up 0.3% (Muck Rack).

So your position isn't what your homepage says. It's what the sources AI trusts say about you. That changes the job. Besides a positioning statement, you need a sentence other people will want to repeat, and the proof that makes it safe to repeat.

How do you find a position worth owning?

Start with the perceptual map, which still works. Plot the attributes that matter most to your buyers on two axes, place each competitor, and look for open ground that buyers care about. Then add the test AI forces on you: can someone else say it back?



The perceptual map still finds open ground. The AI era adds a third test: can others repeat your position?

The three tests of a position

Test	The question	How to check it
Aligned	Does it match what actually motivates buyers to choose?	Win/loss notes, sales calls, the words customers use in reviews and RFPs
Apart	Is it a claim competitors can't make, or won't?	Put your sentence next to the top three competitors' homepages. Could it be theirs?
Repeatable	Can a journalist, a customer or an AI model restate it in one sentence?	Ask five customers what you do. Ask four AI engines. Compare the answers to your sentence.

The repeatable test is the one most teams skip, and the one AI exposes. A position that needs a paragraph of context won't survive a 67-word AI summary. (That's the median length Pew measured for Google's AI summaries; see chapter 5.)

How do you check your position in AI answers?

Run a spot-check. It takes an hour and shows you what buyers see before they ever reach your site.

1. **Write five questions your best buyers ask** before they'd hire someone like you. Use their words, not yours.
2. **Ask each one in four engines**: ChatGPT, Gemini, Perplexity and Google's AI Mode.
3. **Record three things** for each answer: whether you're named, who is named instead, and which sources are cited.
4. **Compare the description to your position.** If AI describes you, does it use your sentence or a vaguer one?
5. **Repeat monthly** with the same questions, so the trend means something.

The spot-check also catches a problem nobody plans for: AI confusing you with someone else. My own firm has this problem. In 16 months of Search Console data, 79% of the search impressions where Google showed us the query were for "PMB" on its own or with other words, and almost none mentioned PMB Direct or Paul. If your name is shared, ambiguous or an acronym, entity clarity is part of positioning.

> **Sidebar: entity clarity in five checks.** - Your company name, the people behind it and what you do read the same way on your site, LinkedIn, directory profiles and author bios. - One plain sentence says what you do, for whom and where. It appears on your About page and in your Organization schema. - Leaders' bios match everywhere, with the same titles, dates and results. - If you share a name or acronym, say what you're not ("not the PMB that…") where it helps. - Third-party profiles (associations, speaker pages, partner sites) use your sentence, not an old one.

What does a strong position sound like?

Specific enough that a competitor couldn't paste it onto their site. The examples below are illustrations, not real companies.

Generic vs. repeatable positioning (illustrative)

AI-average version	Repeatable version
"A trusted partner delivering innovative solutions."	"We rebuild commercial HVAC controls in hospitals without shutting down a single operating room."
"Customer-centric financial advisory for growing businesses."	"We're the CFO bench for family-owned manufacturers in their first generational handoff."
"End-to-end marketing services that drive results."	"We find the spend that isn't working, and prove it before you cut it."

Notice what the right column does. It names who, what and a hard constraint. A reporter can quote it. A model can match it to a question. A buyer can tell a colleague.

What does this look like in practice?

The most durable positioning work I've done was internal before it was external. At Dell Services in 2015, Michael Dell approved the messaging guide, and it was rolled out to the 50,000 services team members worldwide. A position that 50,000 people can say the same way is a position the market hears the same way.

That's the AI-era lesson in one line: consistency across many voices is what makes a position repeatable, and repeatable is what AI picks up.

Checklist: position for the answer

☐ One sentence: what you do, for whom, and the constraint that makes you different
☐ Perceptual map done with buyers' attributes, not yours
☐ Sentence passes the "could a competitor say this?" test
☐ Five-question AI spot-check run in four engines, with results saved
☐ Same sentence on your site, LinkedIn, bios, directory and partner profiles
☐ Proof attached to the position (that's chapter 2)

> **The scoreboard for step 1.** You'll know it's working when (1) AI engines describe you in words close to your sentence, (2) you're named for at least one of your five spot-check questions, and (3) sales hears prospects say your position back to them. Track all three monthly.

Where to go next

Your position is only as strong as the proof behind it. Chapter 2, the proof ladder, shows how to build claims that machines can quote and buyers can check. If you want a senior partner to pressure-test your position with your team, talk to me about advisory support: work with me.

The brand AI names is rarely the loudest one. It's the one everyone describes the same way.

FAQ

What is brand positioning in the age of AI?

It's the place your brand holds in buyers' minds, now shaped largely by what AI models and third-party sources say about you. A strong position is aligned with what motivates buyers, apart from competitors and repeatable in one sentence by someone else.

How do I know if AI describes my brand correctly?

Run a monthly spot-check: ask five real buyer questions in ChatGPT, Gemini, Perplexity and Google AI Mode, and record whether you're named, how you're described and which sources are cited.

Is the perceptual map still useful?

Yes. It's still the fastest way to find open ground buyers care about. Add one more test: whether others can repeat your position in a single sentence.

Why does earned media matter for positioning?

In Muck Rack's 2026 study, 84% of AI citations came from earned media. AI mostly repeats what credible third parties say about you, so your position has to live in their words too.

> Read online: https://www.pmbdirect.com/brand-positioning-alignment-differentiation

STEP 2 OF 7

The Proof Ladder: A Value Proposition and Proof Machines Can Quote

The message ladder still works: a value proposition at the top, pillars in the middle, proof at the bottom. What changed is the bottom rung. Proof now has to be specific enough for a machine to quote and solid enough for a buyer to check.

Step 2 of 7 in The Definitive Guide to Marketing in the Age of AI (in Seven Steps). Originally published October 14, 2024. Rewritten for the AI era October 26, 2026.

A value proposition in the age of AI is only as strong as its weakest proof point. Build it as a **proof ladder**: one value proposition, three supporting reasons, and under each reason a claim that names who, what, how much and when, with a source anyone can check. Adjectives don't survive an AI summary. Specific, dated proof does.

> **In short.** Claim, reason, proof, source, date. If a rung is missing, AI either skips the claim or rewrites it into something vaguer, and buyers trust it less.

Why do generic message pillars fail now?

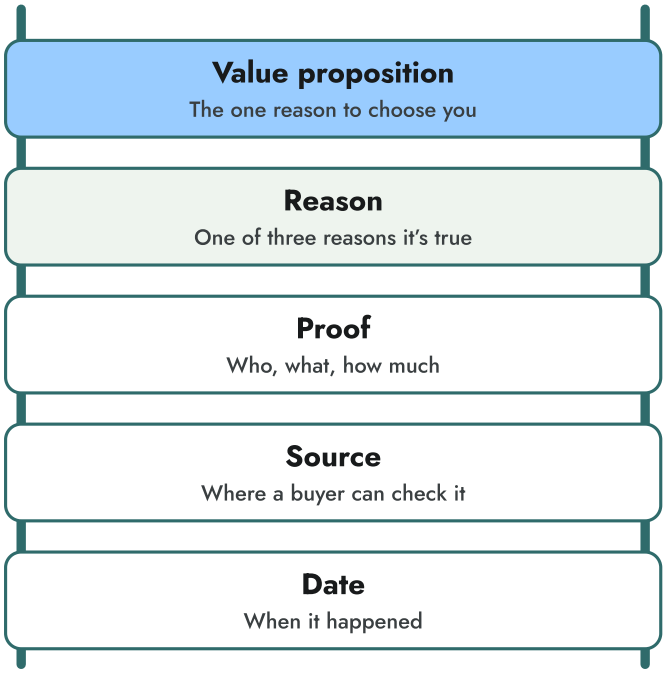
Because they're exactly what AI writes for everyone. "Quality," "innovation," "partnership" and "results" were always weak pillars. In 2026 they're invisible, because a model can produce them for any company in seconds.

The research shows the cost. Among B2B marketers using AI for content, 58% say quality improved and 12% say it went down. Only 39% see better performance (CMI). More words, no more meaning. As CMI's Robert Rose put it, for the most part AI "helps marketers type faster, not think better."

Proof is the one part a model can't supply for you. That makes it the only premium left.

What is a proof ladder?

It's the classic message ladder with the bottom rungs made strict. Each level has one job.



A proof ladder: every claim needs a reason, proof, a source and a date.

The rungs of a proof ladder

Rung	Its job	Weak version	Strong version
Value proposition	The one reason to choose you	"We help companies grow."	"We prove which spend grows revenue, then help you move budget to it."
Reason (pillar)	One of three reasons the proposition is true	"Proven expertise."	"We use mix modeling and experiments, not last-click reports."
Proof	Evidence for the reason	"Clients love our results."	A named result with a number, a scope and a date
Source	Where a buyer can check it	None	A case study, award listing, press article or customer quote with permission
Date	When it happened	None	A year or range, so nobody mistakes old work for new

The examples in the table are illustrations. The structure is the point.

What counts as proof now?

Proof comes in five strengths. Rank yours, and lead with the strongest.

1. **Adjectives** ("best-in-class"). Not proof. Cut them.
2. **Self-reported claims with no number** ("we reduce costs"). Weak. AI often drops them.
3. **Named examples** ("we worked with a national retailer"). Better, but unverifiable.

4. **Numbers with scope and date** (“revenue doubled at D Custom from 2014 to 2024”). Strong. Quotable.
5. **Third-party verification** (an award listing, a press article, a customer quote used with permission). Strongest, because someone else is vouching for you.

Here’s a level-five example from my own work. Texas Heritage for Living, a publication D Custom produced, won Content Marketing Awards from the Content Marketing Institute in 2017, 2020 and 2023. Anyone can look that up on the awards listings. That’s the kind of proof AI can cite and a buyer can verify in a minute.

How do you write a claim a machine can quote?

Write each claim so it stands alone in one sentence. Five rules do most of the work:

1. **One claim per sentence.** Two claims in one sentence get blended or dropped.
2. **Name the entity.** “Dell,” not “a major tech company.”
3. **Give the number its unit and scope.** Revenue for whom? Savings per year, or total? One site, or all of them?
4. **Date it.** A year or a range.
5. **Say whose result it was.** Credit partners. Never take a client’s revenue as your own.

The same claim, before and after

Before	After
“Drove billions in growth for a leading PC brand.”	“At RAPP (1991–2005), the consumer programs I led generated over $1.5 billion in annual revenue for Dell.”
“Award-winning content.”	“D Custom’s client work earned more than 60 national awards while I was president, from 2014 to 2024.”

The “after” column is longer, but every word carries a fact. That’s what survives compression into an AI answer.

Why should you keep a proof bank?

Because proof rots. Titles change, numbers get rounded differently in every deck, and someone eventually writes “billions” where the record says “over $1.5 billion.” A proof bank stops that.

A proof bank is a living record with five columns: the claim, its source, the approved wording, who approved it and when, and every page where it’s used. I keep one for this site, and I checked every result in this guide against it before publishing. When a fact changes, the bank shows every page that needs updating.

> **Sidebar: don't fake the proof.** Buyers notice. In Cint's August 2026 survey of about 1,000 U.S. consumers, AI-generated people were the most disliked kind of AI imagery: 46% said they dislike seeing them in ads (Cint). Show real customers, with permission, or show diagrams. Never a synthetic "customer."

Checklist: build your proof ladder

☐ One value proposition, in one sentence a buyer could repeat
☐ Three reasons, each a claim a competitor couldn't copy
☐ At least two proof points per reason, ranked by the five levels above
☐ Every proof point has a number with scope, a date and a source
☐ Partners credited; client revenue described as the client's
☐ A proof bank with approved wording, owner and every page where each claim is used

> **The scoreboard for step 2.** Count the share of claims on your top ten pages that have a number, a date and a source. Aim for all of them. Then track sales use: how often reps send or quote a proof point in a deal. Proof that sales doesn't use isn't doing its job.

Where to go next

Once you know what you can prove, plan where to say it. Chapter 3, the question map, turns your proof into a plan built on the questions buyers and their agents actually ask. If you want help building the ladder and the proof bank with your team, talk to me about advisory support: work with me.

Anyone can claim. Only you can prove it, and only if you wrote it down.

FAQ

What is a proof ladder?

A message ladder with strict bottom rungs: a value proposition, three supporting reasons, and under each reason proof that names who, what, how much and when, with a checkable source.

What's the difference between a value proposition and a message pillar?

The value proposition is the single reason to choose you. Message pillars are the two or three reasons it's true. Proof points are the evidence for each pillar.

How do I make proof points AI can cite?

Write one claim per sentence, name the entity, give the number its unit and scope, add a date and link a source a buyer can check.

What is a proof bank?

A living record of every claim you make, with its source, approved wording, who approved it and when, and where it’s used, so facts stay consistent everywhere.

> Read online: https://www.pmbdirect.com/the-message-ladder

STEP 3 OF 7

The Question Map: Plan Around What Buyers and Their Agents Ask

The old editorial calendar planned topics by date and volume. In the age of AI, plan by question: which questions your buyers and their AI agents ask, and which single page of yours answers each one.

Step 3 of 7 in The Definitive Guide to Marketing in the Age of AI (in Seven Steps). Originally published October 14, 2024. Rewritten for the AI era October 26, 2026.

An editorial calendar for the age of AI starts with questions, not dates. List the questions your buyers ask at each stage, and note whether a person, an AI assistant or an agent is asking. Give each question one owner page on your site, plus the outside sources you want answering it too. Then publish fewer, better pieces against that map.

> **In short.** Questions, not topics. One owner page per question. And plan for the outside sources AI will quote, not just your own pages.

A note on proof: Proof below predates generative AI; an AI-era result replaces it by mid-December.

Why does a calendar built on volume fail now?

Because volume is a commodity and the reader has changed. In HubSpot's 2026 survey, 83.5% of marketers said they're expected to produce more content (HubSpot). Every competitor can hit any publishing quota with AI. Nobody wins on quota.

Meanwhile, questions are exactly what AI answers first. Pew found that 60% of Google searches starting with "who," "what," "when" or "why" produced an AI summary (Pew). Seer Interactive found question-style queries triggered an AI Overview 85.9% of the time (Seer Interactive). If your buyers' questions get answered before the click, your plan has to be built around being the source of those answers.

The old version of this guide proves the point. It was a calendar-driven series, and in 16 months the whole series earned 89 search impressions and no clicks. It answered how to run a content program. Almost nobody was asking that.

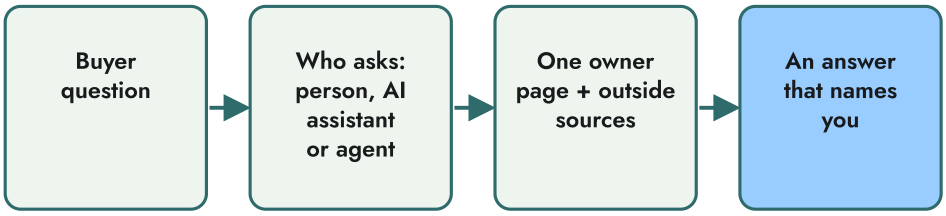
What is a question map?

It's the persona-by-stage content matrix from the old calendar, turned on its side. Rows are buyer questions. Columns tell you who asks, where the answer should live and what proof it needs.

A question map (template with example rows)

Buyer question	Stage	Who asks	Owner page	Outside sources to earn	Proof needed
"How do we measure marketing without clicks?"	Problem	CMO, via AI assistant	Your measurement guide	Trade-press byline, podcast	A dated result with a number
"Who does this well for companies like ours?"	Shortlist	Buying group, via AI and agents	Your industry page	Analyst or association mention, client quote	Named clients in that industry
"What does it take to switch?"	Decision	Champion, then a rep	Your process page	Customer review sites	Timeline, team, risks addressed

The example rows are illustrations. Build yours from real questions.



One question, one owner page. Plan the earned sources too.

A question map plans by question, owner page and outside sources.

Where do the real questions come from?

From people who talk to buyers, and from the machines buyers talk to. Gartner's survey of 645 B2B buyers found they used an average of seven information sources in a recent purchase (Gartner). Mine as many of those as you can:

- **Sales calls and win/loss notes.** The questions that come up right before a deal stalls are the ones to own.
- **Customer support and onboarding.** Post-sale questions are tomorrow's pre-sale questions.
- **Search Console queries.** Long, question-shaped queries are a direct line to what people type.
- **AI spot-checks.** Ask your buyers' questions in four engines and note what's missing or wrong.
- **RFPs and procurement forms.** They tell you what a buying group must document.

Why does each question need one owner page?

Because AI quotes one passage, and Google picks one page. Two of your pages competing for the same question split whatever authority you have, and often the wrong one wins.

Give every question a single canonical page. Other pages can mention the topic, but they link to the owner instead of trying to rank for it. When you add a new piece, check the map first. If a page already owns that question, improve it instead of writing another one.

How do you plan for agents, not just people?

Assume part of the research will be done by software acting for the buyer. Google announced information agents that monitor the web for a user, and agentic booking for local services, in May 2026 (Google). An agent building a shortlist doesn't read your story. It looks for facts it can match to its instructions.

So put the facts an agent needs on every shortlist-stage owner page, in plain text. Say what you do, for whom, where and how you work. Say what it takes to start and what proof you have. If an agent can't find a fact, it can't put you on the list.

> **Sidebar: plan the earned column, not just the owned one.** AI answers mostly cite other people's content (chapter 1 has the numbers). For your top questions, plan at least one outside source a month: a guest article, a podcast, an association talk or a partner write-up. Chapter 6 covers how.

What cadence works?

A rhythm tied to questions and proof, not to the calendar. One strong owner page a week or two, each improving or replacing a weaker page. One earned placement a month. A monthly spot-check to see which questions you now own. Fewer pieces, each with a job.

Checklist: build your question map

☐ 20–40 real buyer questions, collected from sales, support, search data and AI spot-checks
☐ Each tagged by stage and by who asks (person, AI assistant, agent)
☐ One owner page per question, with no two pages competing
☐ Agent-ready facts on every shortlist-stage page
☐ An outside source planned for each of your top ten questions
☐ A cadence you can sustain with quality, not a quota

> **The scoreboard for step 3.** Track, per question: whether your owner page shows up in search impressions, whether AI engines name you for it, and whether sales hears it from prospects. A question you own shows up in all three.

Where to go next

Owning the question is half the job. The page also has to do the work the buyer came for. Chapter 4 covers content that does the buyer's work, and their agent's. If you want an outside read on which questions AI already answers for your buyers, and whether it names you, ask about the measurement and AI-visibility audit: work with me.

The calendar asks, "What should we publish next?" The map asks, "What will our buyer ask next, and who will answer?"

FAQ

What is a question map in content marketing?

A content plan built on buyer questions instead of topics and dates. Each question is tagged by buying stage and by who asks (a person, an AI assistant or an agent), and is assigned one owner page plus outside sources to earn.

Do I still need an editorial calendar?

Yes, as a schedule. But plan the calendar from the question map, so every piece answers a question you've chosen to own instead of filling a quota.

How do AI agents change content planning?

Agents look for facts they can match to instructions. Put plain-text facts on shortlist-stage pages: what you do, for whom, where, how you work, what it takes to start, and your proof.

How many pages should answer one question?

One. Other pages can mention the topic and link to the owner page, but only one page should target each question.

> Read online: https://www.pmbdirect.com/the-editorial-calendar

STEP 4 OF 7

Do the Buyer's Work, and Their Agent's: How to Win on Value

Great content is still a product. In the age of AI, it has two customers: the buyer, and the AI the buyer sends to do the research first.

Step 4 of 7 in The Definitive Guide to Marketing in the Age of AI (in Seven Steps). Originally published November 4, 2024. Rewritten for the AI era October 26, 2026.

Content wins on value when it does work the buyer would otherwise have to do: frame the problem, compare the options honestly, build the business case and lower the risk of saying yes. Now that buyers send AI to do the first pass, that work has to be readable by a machine and checkable by a person. Do both and you compete on value. Skip them and you compete on price.

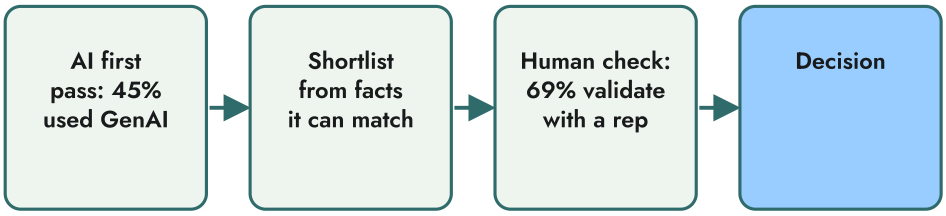
> **In short.** Buyers delegate research to AI, then check it with people. Give the machine facts it can match and give the champion proof they can defend.

A note on proof: Proof below predates generative AI; an AI-era result replaces it by mid-December.

How are B2B buyers using AI to buy?

To gather information first, and to decide later with people. Gartner's survey of 645 B2B buyers, fielded in August and September 2025, found that 45% used generative AI in a recent purchase, mostly to gather information on vendors and products. A later Gartner presentation put GenAI use during purchases at 60% (Gartner; Gartner).

But buyers don't fully trust the machine. In the same survey, 69% said they prefer to validate AI-generated insights with a sales rep. And 51% said they're more likely to meet misleading information from GenAI (49% said the same of reps). So the journey now has two checkpoints: what the AI found, and what a person confirms.



Source: Gartner survey of 645 B2B buyers, Aug–Sep 2025.

Buyers let AI do the first pass, then check it with people (Gartner, 2026).

What work is the buyer trying to get done?

Five jobs, roughly in order. Your content should do each one better than the buyer could alone.

The buyer's five jobs, and what content does for each

The buyer's job	Content that does it	What an AI agent needs from it
Understand the problem	A plain explanation of causes and costs, with numbers	Definitions and figures stated in one sentence each
Compare the options	An honest comparison, including when you're not the right fit	A table with the same criteria for every option
Build the case	A business-case outline with the inputs a CFO will ask for	The inputs listed plainly: costs, timeline, team, risks
Lower the risk	Proof from similar companies, implementation steps, what can go wrong	Named, dated results; a step list; stated constraints
Get sign-off	A short summary the champion can forward	A standalone summary paragraph at the top

The last row matters most. Your real reader is often the champion who has to defend the decision in a meeting you'll never attend. Arm that person.

Why do buyers still call a rep, and what should the rep carry?

Because AI can't take responsibility for a decision, and a buying group needs someone who can. That 69% is good news for sellers who show up with proof. It's bad news for those who show up with a pitch.

Build validation assets for reps:

- **A one-page "what AI gets wrong about us"** that corrects common errors with sources.
- **The comparison table** from your site, so the rep and the AI tell the same story.
- **Three dated, named results** the rep can cite from memory.
- **A risk list**: what can go wrong, and how you handle it.

How does content defend price?

By making value easier to see than cost. When buyers can't tell options apart, price is the only thing left to compare.

I've seen this work. In 2020, during my D Custom years, McKesson's new pitch reduced price pressure and increased its win rate. At D Custom (2014–2024), thought leadership written for generalists rather than specialists helped DDN shorten purchase cycles and win more deals.

What makes content agent-readable?

Clear facts, stated once, in the same words everywhere. An agent matching your company against a buyer's instructions needs:

1. **A one-sentence description** of what you do, for whom and where.
2. **Structured comparisons** in HTML tables, not images of tables.
3. **Constraints stated plainly**, such as minimum project size, regions served or what you don't do.
4. **Proof with names and dates**, linked to a page that holds the detail.
5. **Consistent entity details**, so the agent doesn't confuse you with someone else (see chapter 1).

This also helps industry pages. Specific beats general for both people and machines, which is why I recommend publishing industry-specific content.

> **Sidebar: say when you're not the right fit.** It feels risky. It works. An honest "not for you if…" line saves the buyer time, earns trust with the champion and gives an agent a clean reason to include you when the fit is right. The agent will find the limits anyway. Better that it finds them in your words.

Checklist: do the buyer's work

☐ One page per buyer job, or one page that does all five clearly
☐ An honest comparison table, with the same criteria for every option
☐ A business-case outline with the inputs a CFO asks for
☐ A forwardable summary paragraph at the top of every key page
☐ Validation assets for reps, including "what AI gets wrong about us"
☐ Plain-text facts and constraints an agent can match

> **The scoreboard for step 4.** Watch for fewer price objections, shorter time from first meeting to decision, and how often reps and champions forward your pages. Ask new clients which page they sent to their boss. That page is doing the work.

Where to go next

Doing the buyer's work takes clear writing that people and machines can both use. Chapter 5 covers how to write for the skimmer, the scanner and the machine, and how an AI-native team produces it. If you want to see what AI tells your buyers about you today, and what your content is worth to the pipeline, ask about the measurement and AI-visibility audit: work with me.

Buyers don't pay more for nicer words. They pay more when you've already done half their job.

FAQ

What does "win on value" mean in B2B marketing?

It means giving buyers enough clarity about outcomes, fit and risk that price stops being the deciding factor. Content does that by doing the buyer's research and business-case work for them.

How do B2B buyers use AI in purchasing?

Gartner found 45% of B2B buyers used generative AI in a recent purchase, mainly to gather vendor and product information, and 69% prefer to validate AI-generated insights with a sales rep.

What content should sales reps have in the AI era?

Validation assets: a page correcting what AI gets wrong about you, a comparison table that matches your site, three dated and named results, and a plain risk list.

How do I make my content useful to AI agents?

State facts plainly and consistently: what you do, for whom and where, your constraints, structured comparisons in real tables and named, dated proof.

> Read online: https://www.pmbdirect.com/content-how-to-win-on-value

STEP 5 OF 7

Write for the Skimmer, the Scanner and the Machine

People still decide on feeling and justify it with reasons. That hasn't changed. What changed is that a machine now reads first, and passes along only the passages it can lift whole.

Step 5 of 7 in The Definitive Guide to Marketing in the Age of AI (in Seven Steps). Originally published November 19, 2024. Rewritten for the AI era October 26, 2026.

Writing in the age of AI means writing for three readers at once: the **skimmer** who reads only headings, the **scanner** who reads tables and callouts, and the **machine** that lifts single passages into an answer. Lead every section with the answer, make every passage stand alone, and keep the persuasion principles that always worked. Then produce it with an AI-native team that has a brief, a fact check and a human who approves.

> **In short.** Answer first. One idea per passage. Proof you can check. And staff AI like a team, not a typewriter.

A note on proof: Proof below predates generative AI; an AI-era result replaces it by mid-December.

Who reads your content now?

Three readers, in this order. The machine often reads first, the skimmer decides whether to stay, and the scanner decides whether to believe you.

The three readers and what each one needs

Reader	What they read	What they need from you
The machine	Single passages, often only the first lines of a section	A direct answer in the first two sentences, with named entities and numbers
The skimmer	Title, headings, bold text	Headings that are real questions or clear claims, never teasers
The scanner	Tables, callouts, lists, captions	Comparisons and checklists that stand on their own

Here's the number that should change how you write. In Pew's study of real Google searches, the median AI summary was 67 words long (Pew). Whatever you want buyers to take away has to survive being squeezed into about one short paragraph. So write that paragraph first.

How do you write a passage AI can lift?

Treat every section like a standalone answer. Five habits do most of the work:

1. **Lead with the answer.** The first two sentences under a heading should answer the heading.
2. **One idea per passage.** If a paragraph makes two points, split it.
3. **Name things.** Say “Gartner’s 2026 CMO Spend Survey,” not “a recent study.”
4. **Put numbers next to their scope.** “45% of 645 B2B buyers surveyed in 2025,” not “nearly half of buyers.”
5. **Use real HTML tables and lists** for comparisons and steps, so both scanners and machines can read them.

None of this makes writing dull. A clear answer up front earns the right to tell the story underneath it.

Do classic persuasion principles still work?

Yes. People haven’t changed. The principles that made the original version of this chapter useful still hold, but each needs an honest AI-era check.

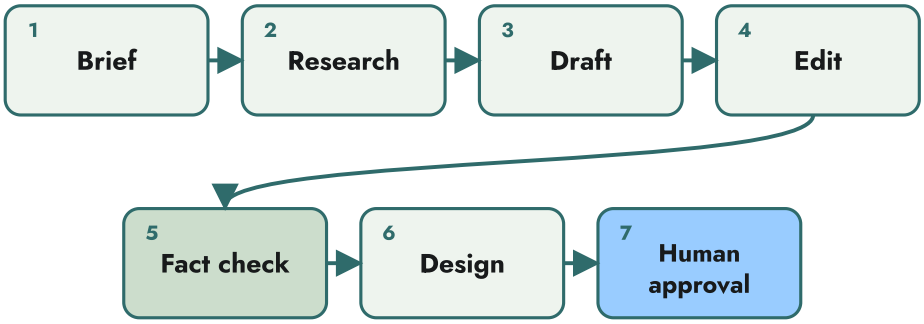
Persuasion principles, with an AI-era check

Principle	What it means	AI-era check
Priming	What readers see first shapes how they read the rest	The first lines may be all the machine passes on. Make them your point.
Social proof	People follow people like them	Use only real, named, dated proof. Never a synthetic customer.
Clustering	Numbered lists and clear headings help memory	They also help machines parse your structure.
Verbatim effect	People remember the gist, not the words	Make the gist specific. “Over $1.5 billion in annual revenue for Dell” sticks. “Billions” doesn’t.
Story	Stories connect, inspire and move people	Keep one specific scene per piece. AI can’t invent your experience.
Scarcity	Limited supply raises value	Only in offers, and only when it’s true. Never in articles.

The underlying idea is the same one the old chapter started with: we decide on feeling and justify the decision with reasons afterward. Earn the feeling, then supply the proof. More on that in Why We Like Stories.

How should an AI-native team produce content?

Like a content team, not a typewriter. The mistake is asking AI for “a draft” and editing whatever comes back. The fix is to give AI the roles a good editorial team has, with gates between them. I described the method in Stop Asking AI for a Draft. Staff It Like a Content Team.



A gate between every step. A named person approves before anything goes live.

Staff AI like a content team, with a gate between every step.

1. **Brief.** One page: the reader, the single takeaway, the proof you'll use and the action you want.
2. **Research.** Sources opened and logged, with dates. Nothing invented.
3. **Draft.** Answer-first, against the brief.
4. **Edit.** Cut 15–25%. Check headings, rhythm and the one call to action.
5. **Independent fact check.** A fresh pass that didn't write the piece opens every source and checks every claim against the proof bank.
6. **Design.** Diagrams and tables that teach, with alt text.
7. **Human approval.** A named person signs off before anything goes live.

Most teams aren't set up for this yet. In Gartner's 2026 CMO Spend Survey, 70% of CMOs said their internal processes aren't mature enough to implement and scale AI, and only 30% reported mature AI readiness (Gartner). The gap is process, not tools.

What does AI expose about your team?

Everything your team never wrote down. AI is only as good as the context you give it: the positioning decisions, the approved claims, the reasons a past campaign was killed. In most teams, that lives in people's heads and old decks. When those people leave, the context goes with them, and AI fills the gap with the average of the internet.

So the first AI investment isn't a tool. It's a written record of what you know and why you decided it. I make the full case in AI Won't Replace Your Marketing Team. I have a stake in it: I'm Chief Marketing Officer of Nvitis, an institutional memory and context layer. But you don't need any product to start. A shared proof bank and a decision log will do (see chapter 2).

> **Sidebar: tell people how you use AI.** In Cint's August 2026 survey of about 1,000 U.S. consumers, 74% said brands have a duty to disclose AI use in advertising, and 44% said finding undisclosed AI ads would lower their trust (Cint). Write a one-paragraph policy: where you use AI, where a person always decides and what you'll never generate (for example, images of real-looking people).

What still takes a human?

Judgment, taste and experience. Design is a good example. Texas Heritage for Living, which D Custom produced, won three rounds of Content Marketing Awards: Best Use of Photography and Best Overall Design (Print) in 2017, Best Overall Design (Print) in 2020 and Best Print Publication in 2023. Those wins came from people making hundreds of small choices about what to leave out. AI can speed up the making. It can't decide what's worth making.

Checklist: write for three readers

☐ Every section answers its heading in the first two sentences
☐ Headings are real questions or clear claims
☐ One idea per passage; numbers next to their scope and source
☐ At least one real table and one standalone callout per piece
☐ A brief, an independent fact check and a named approver for every piece
☐ A written AI-use policy, published where readers can find it

> **The scoreboard for step 5.** Track cycle time from brief to approval, the number of claims caught at fact check, and whether AI answers quote your passages for your target questions. If output goes up and none of those move, you're typing faster, not thinking better.

Where to go next

Well-written content still needs to reach the places AI reads. Chapter 6 covers how to earn citations instead of chasing clicks. If you want to set up an AI content team with real gates, talk to me about building your AI content team: work with me.

AI can write the sentence. Only your team knows which sentence is true.

FAQ

How do I write content that AI search will quote?

Answer each heading in the first two sentences, keep one idea per passage, name entities, put numbers next to their scope and source, and use real HTML tables and lists.

How long is a typical AI summary?

In Pew Research Center's study of real Google searches from March 2025, the median AI summary was 67 words. Write your key takeaway so it fits in about one short paragraph.

What is an AI content team?

A workflow that gives AI the roles of an editorial team (research, drafting, editing, fact checking, design) with gates between them and a human who approves before publishing.

Should brands disclose AI use?

Cint's 2026 survey found 74% of U.S. consumers think brands should disclose AI use in advertising. A short, public AI-use policy is a low-cost way to protect trust.

> Read online: https://www.pmbdirect.com/writing-and-design

STEP 6 OF 7

From Traffic to Citations: How to Get Named in AI Answers

Content amplification used to mean driving traffic back to your website. Now the answer often appears before the click, so the goal is to be the source that answer cites, and the brand it names.

Step 6 of 7 in The Definitive Guide to Marketing in the Age of AI (in Seven Steps). Originally published December 16, 2024. Rewritten for the AI era October 26, 2026.

To get cited in AI answers, be where AI reads. That takes three channels working together. **Earned sources**: journalism, guest articles, podcasts, associations and partner sites. **Owned pages** built so a single passage can be quoted. **Paid and in-person channels** that still reach buyers. Measure success by whether you're named and cited for your buyers' questions, not by traffic alone.

> **In short.** Traffic is the wrong scoreboard. Being named, cited and chosen is the right one. Most AI citations come from what other people publish about you.

A note on proof: Proof below predates generative AI; an AI-era result replaces it by mid-December.

Is website traffic still the right goal?

Not on its own. When Google shows an AI summary, people click a regular search result on 8% of visits, compared with 15% when there's no summary. They click a link inside the summary on just 1% of visits, and they end their browsing session more often (26% vs. 16%) (Pew). Ahrefs found AI Overviews go with a 58% lower click-through rate for the top-ranked page (Ahrefs).

The buyer still learns something. They just learn it from the answer. If the answer names you, you've won the moment even without a visit.

Does being cited in an AI answer actually help?

Yes, though it doesn't bring back the old traffic. Seer Interactive's 2026 study of 53 brands found that being cited in an AI Overview delivered about 120% more organic clicks per impression than not being cited on the same results page. It still delivered 38% fewer than pages with no AI Overview at all (Seer Interactive). Seer is careful to call this correlation, not proof of cause.

So citations are the new rank. They're the best position left on an AI results page.

Where do AI answers get their sources?

Mostly from people other than you. Muck Rack's May 2026 study of more than 25 million links cited by ChatGPT, Claude and Gemini found:

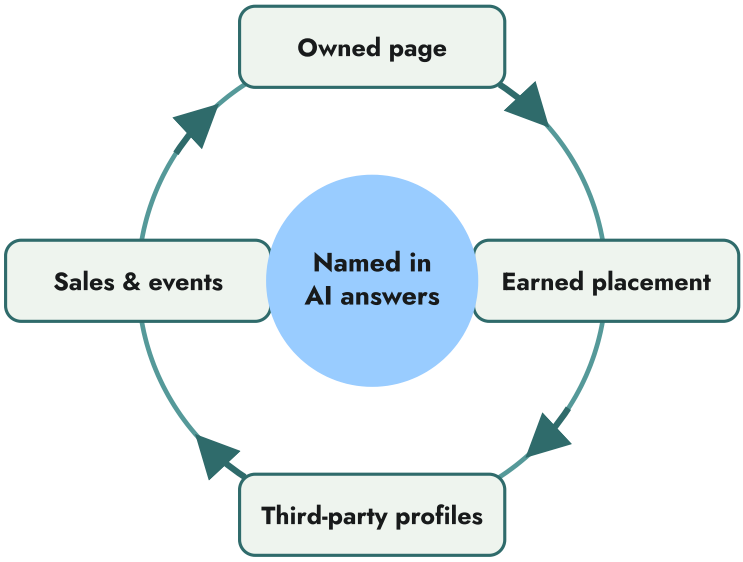
- **84%** of citations came from earned media, a share that has held between 82% and 89% since July 2025.
- **27%** came from journalism alone.
- **0.3%** came from paid and advertorial content.
- More than half of journalism citations came from articles published in the past 12 months.

Each engine behaves differently. ChatGPT cited sources in 96% of responses, averaging five. Gemini cited in 82%, averaging eight. Claude cited in 55%, averaging 13 (Muck Rack).

The takeaway is uncomfortable for content teams: your blog alone won't get you named. It needs outside sources pointing to the same position and proof.

What is the citation flywheel?

It's how owned, earned and in-person channels feed each other. Each turn makes the next one easier.



The citation flywheel: owned, earned and in-person channels feed each other.

1. **Owned page.** One canonical page per buyer question, with an answer-first passage and dated proof (chapters 3 and 5).
2. **Earned placement.** A guest article, podcast, association talk or press quote that repeats your position and links to the owned page.

3. **Third-party profiles.** Speaker pages, partner listings and bios that describe you the same way.
4. **Sales and events.** Reps and speakers use the same proof, which produces new stories, quotes and invitations.
5. **Back to owned.** Each new result updates the owned page, so it stays recent.

Speaking is a good example of the earned turn. I've given more than 10 keynotes, at SMU, the University of Michigan and the University of Washington. A keynote reaches the room once. The host's write-up, the recording page and the speaker bio keep describing you long after.

What are the steps to get cited in AI answers?

Generative engine optimization (GEO) sounds technical. Most of it is disciplined marketing. Here are the steps in order:

1. **Pick 5–10 buyer questions you want to own** (from your question map).
2. **Run a baseline spot-check** in ChatGPT, Gemini, Perplexity and Google AI Mode. Record who's named and which sources are cited.
3. **Fix the owned page** for each question: answer first, one idea per passage, named and dated proof, FAQ and schema.
4. **Target the sources the engines already cite.** If one trade publication or association shows up again and again, that's where to pitch.
5. **Earn one placement a month** that carries your position and proof.
6. **Make your entity consistent** across your site, LinkedIn, bios and directories.
7. **Re-run the spot-check monthly** with the same questions, and log changes.

> **Sidebar: what should paid media do now?** Keep it for what it does reliably. Seer found paid search held up best of all search channels through 2025 and early 2026, with click-through rates on AI Overview pages steady at 13% to 16%. Seer calls it "your most predictable lever in search right now." Use paid to cover high-intent queries and to promote earned placements, and judge it on cost per qualified conversation, not clicks.

What still works from the old amplification playbook?

Organic and paid working together, events and sales as channels, and repetition. People still need to see an idea several times before it sticks. The difference is the destination. The original guide said every channel should point back to your website. Now every channel should point to the same position and proof, wherever the buyer meets it: in an AI answer, a podcast, a rep's email or your site.

Checklist: from traffic to citations

☐ 5–10 target questions, with a baseline spot-check saved

☐ An answer-first owned page for each
☐ A list of the outside sources AI already cites for those questions
☐ One earned placement a month, carrying your position and proof
☐ Consistent descriptions on every profile and bio
☐ Paid judged on qualified conversations, not clicks

> **The scoreboard for step 6.** Monthly: in how many of your target questions are you named, in how many engines, and which of your earned placements show up among the citations? Alongside that, track search impressions (not just clicks) for your owned pages, and ask every new enquiry where they first heard of you.

Where to go next

Citations and mentions are leading indicators. The real question is what grows revenue. Chapter 7 covers how to measure that when clicks disappear. If you want a baseline of where AI answers name you today, and what to fix first, ask about the measurement and AI-visibility audit: work with me.

You can't buy your way into most AI answers. You have to be worth quoting, by people other than you.

FAQ

How do I get my company cited in AI answers?

Own a few buyer questions with answer-first pages, earn placements in the outside sources AI already cites for those questions, keep your company description consistent everywhere, and re-check monthly in four engines.

What is generative engine optimization (GEO)?

The practice of making your brand more likely to be named and cited in AI-generated answers. Most of it is clear positioning, checkable proof, answer-first pages and earned media.

Do AI Overviews reduce website traffic?

Studies say yes. Pew found people clicked a regular result on 8% of visits with an AI summary vs. 15% without, and Ahrefs measured a 58% lower click-through rate for the top-ranked page.

Does earned media matter for AI search?

Yes. Muck Rack's May 2026 study found 84% of AI citations came from earned media and 27% from journalism, while paid and advertorial content was 0.3%.

> Read online: https://www.pmbdirect.com/content-amplification

STEP 7 OF 7

Measure What Grows Revenue When Clicks Disappear

A headline test can't tell you where to move a budget, and last-click tracking can't see a sale that started inside an AI answer. Here's what can.

Step 7 of 7 in The Definitive Guide to Marketing in the Age of AI (in Seven Steps). Originally published February 24, 2025. Rewritten for the AI era October 26, 2026.

To measure marketing when clicks disappear, use four instruments, each for its own job. **Mix modeling** decides how to split the budget. **Incrementality experiments** prove what a channel adds. **A/B tests** settle smaller choices on pages and messages. **AI-visibility checks** give you a leading indicator. Then write the decision rule before you read the result. An insight without a decision rule is just another report.

> **In short.** Models for the budget, experiments for the truth, A/B tests for the details, spot-checks for AI. And decide in advance what you'll do with each answer.

Why does click-based attribution break in the age of AI?

Because the click is no longer the first thing that happens. A buyer asks an AI assistant and reads a summary that names three vendors. She hears your name again on a podcast. Two weeks later, she types your company name into a browser. Last-click attribution gives all the credit to that last search. Everything that made it happen is invisible.

When Google shows an AI summary, people click a regular result about half as often (Pew; chapter 6 has the numbers). Attribution has less to count, while buyers keep learning from the answer.

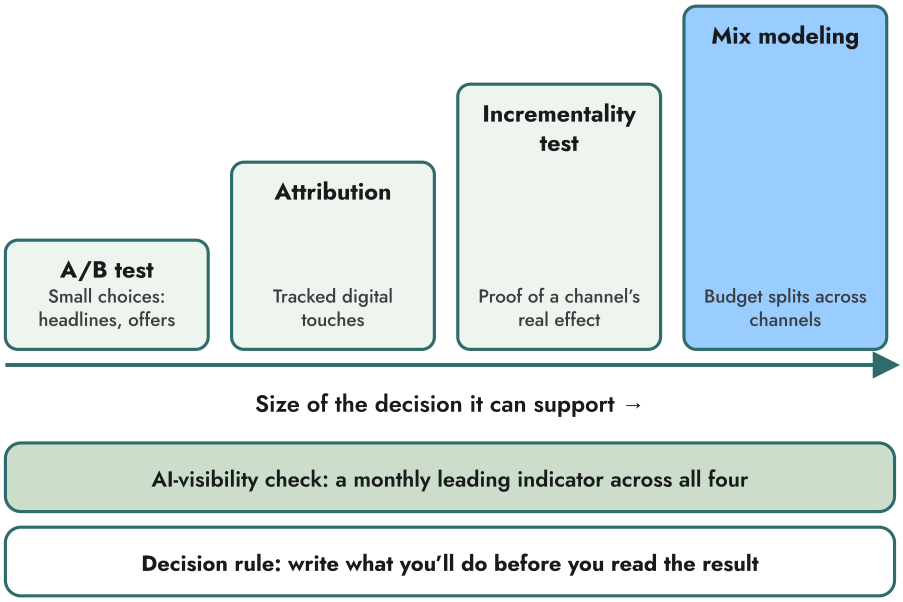
Marketers feel it. Measuring marketing ROI was the top challenge in HubSpot's 2026 survey, cited by 33% (HubSpot). In CMI's research, 33% of B2B marketers named measuring content effectiveness as a top challenge (CMI). And with budgets flat, you have to defend every dollar.

What's the difference between mix modeling, incrementality and attribution?

They answer different questions. Use each for its own job.

Four measurement instruments and when to use each

Instrument	The question it answers	Best for	Watch out for
Marketing mix modeling (MMM)	How much does each channel contribute to revenue, and where should the next dollar go?	Budget splits across channels, including ones with no clicks	Needs enough history and spend variation; answers come in weeks, not days
Incrementality experiment	What happens if we turn this channel off, or up, for a test group?	Proving a channel's real effect before a big budget move	Needs a clean test and control design
A/B test	Which version of a page, email or message works better?	Headlines, offers, formats, calls to action	Tells you about small choices, not budget splits
Multi-touch or last-click attribution	Which tracked touches came before a conversion?	Day-to-day optimization inside digital channels	Blind to AI answers, word of mouth and offline influence



Match the instrument to the decision, and write the decision rule first.

What did mix modeling find at a single Harrah's property?

That more than $30 million a year in promotional spend could go without hurting results. In 2012, working in partnership with Polaris Research, we used marketing mix modeling to optimize the promotional campaigns at a single Harrah's property. The work cut more than $30 million a year from marketing spend, with performance maintained.

The lesson isn't the size of the number. It's the order of operations: model first, then cut. Let the evidence show what is and isn't contributing, and let the budget follow it. AI makes that habit worth more, because AI makes it easy to fund whatever is easiest to count.

Disclosure: Paul is EVP, Business Development at Polaris Research, the modeling partner on the Harrah's work.

Why do so few companies act on mix modeling?

Because they treat it as a report instead of a decision tool. In an HBR survey sponsored by Google, 73% of respondents said they're under growing pressure to show what their media channels do. Only 22% were "leaders," very effective at using mix modeling to act on insights. Among leaders, 55% connect model outputs to real marketing decisions, compared with 20% of laggards (EMARKETER).

The gap is decisions, not data. Which brings us to the most useful habit in this guide.

What is a decision rule, and why write it first?

A decision rule says what you'll do with a result before you see it. It turns measurement into action and stops anyone from explaining away an inconvenient answer.

Decision rules (template with example rows)

If the read shows…	Then we will…	Owner	By when
A channel's modeled return is below our threshold for two reads in a row	Cut its budget by a set share and run an incrementality test	CMO	Next budget cycle
A test group with the channel off shows no drop in qualified conversations	Move that budget to the best-performing channel	Head of growth	Within 30 days
We're named in fewer than one of five target questions after 90 days	Shift effort from owned pages to earned placements	Content lead	Next quarter

The rows are illustrations. Your thresholds and owners are yours to set. The point is that they're written down before the data arrives.

Is A/B testing still worth doing?

Yes, for the decisions it fits. Test headlines, offers, page structure and calls to action. Test one thing at a time. Decide in advance what "better" means. Run it long enough to trust. What A/B testing can't do is tell you whether a channel deserves its budget. That's a job for models and experiments.

Models can make even small choices pay off. For Cross, in work through PMB from 2010 to 2017, a predictive targeting model more than doubled the revenue from each campaign. At RAPP (1991–2005), the consumer programs I led generated over $1.5 billion in annual revenue for Dell. We also used predictive modeling there to cut acquisition costs for clients such as AT&T and American Express. Testing tells you which message wins. Modeling tells you who to send it to.

How do you measure AI visibility when there's no click?

With a fixed set of questions, checked the same way every month. AI visibility is a leading indicator. It moves before revenue does.

1. **Pick five questions** your buyers ask before they'd hire you, and keep them fixed.
2. **Check four engines monthly**: ChatGPT, Gemini, Perplexity and Google AI Mode.
3. **Log three things**: whether you're named, who's named instead and which sources are cited.
4. **Watch impressions and clicks separately** in Search Console. Seer Interactive found cases where impressions more than doubled while clicks held flat, so click-through rate alone can mislead (Seer Interactive).
5. **Ask every enquiry** where they first heard of you, and record answers that mention AI tools.

> **Sidebar: read the scoreboard in this order.** 1. Qualified conversations that fit an offer, by source 2. Sales use: how often reps send or cite a piece 3. Search impressions for target questions 4. AI-answer mentions and citations 5. Traffic, only to diagnose problems
>
> Revenue signals first, attention signals last.

Checklist: measure what grows revenue

☐ Each budget question matched to the right instrument (model, experiment, A/B test or attribution)
☐ Mix modeling or a lighter-weight model in place for channel splits
☐ At least one incrementality test planned for your biggest channel
☐ Decision rules written before each read, with owners and dates
☐ A fixed five-question AI spot-check, run monthly in four engines
☐ An enquiry form or intake call that asks where people first heard of you

> **The scoreboard for step 7.** The proof that measurement works is a budget decision you made because of it. Count them. If last quarter's reads didn't change a single allocation, you have reporting, not measurement.

Where to go next

Measurement closes the loop and starts the next one. Take what you learned back to chapter 1. For the same open loop in content, read Thought Leadership and Curriculum Fail the Same Way. Want an outside read on which spend grows revenue and whether AI answers name you? Ask about the measurement and AI-visibility audit: work with me.

The click was always a proxy. Now that it's fading, measure the thing it stood for.

FAQ

How do you measure marketing without clicks?

Use mix modeling to split the budget, incrementality experiments to prove what each channel adds, A/B tests for smaller choices and a monthly AI-answer spot-check as a leading indicator. Write decision rules before you read the results.

What is marketing mix modeling?

A statistical model that estimates how much each marketing channel contributes to sales or revenue, including channels without clicks, so you can decide where the next dollar should go.

What's the difference between MMM and attribution?

Attribution counts tracked touches before a conversion and misses AI answers, word of mouth and offline influence. Mix modeling estimates each channel's contribution from spend and outcome data over time.

What is an incrementality test?

An experiment that turns a channel off or up for a test group and compares results with a control group, to show what that channel really adds.

Is A/B testing still useful?

Yes, for headlines, offers, formats and calls to action. It can't tell you how to split a budget across channels; that takes models and experiments.

> Read online: https://www.pmbdirect.com/testing-dont-skip-this-step