

# State of Content Marketing and AEO in an AI-Driven Buying Era

## Executive Summary

*(Please note: I like, enjoy, and favor the use of the “—” dash as a writer and editor and content marketer of fifteen years. Did I use AI to make this? Sure for some aspects. But you’ll agree with me on “why” and “how” by the end of this article).*

In the last five years, I’ve watched a painfully obvious disconnect grow between customers seeking a solution to their problems and brands offering that Cinderella “fit” glass slipper. The problem pushing them apart? The providers aren’t actually listening to what **potential buyers or even customers want or how they consume content**. So we now have a market where brands waste time pushing stale marketing tactics that didn’t even work five years ago and wonder why they only hear crickets.

If you experience how B2B buyers actually research and evaluate software over the last 12–18 months, the pattern is hard to ignore. Research doesn’t start the way it used to. More and more, it starts with a prompt and an AI/LLM solution—and before a buyer ever visits your site, they may already have a shortlist and a point of view shaped entirely by what an AI system told them.

The data has caught up to what a lot of practitioners have been observing in the field. Gartner’s 2026 sales survey found that 67% of B2B buyers now prefer a rep-free experience, and 45% actively used AI during a recent purchase—not as a novelty, but as a functional research tool. G2’s 2025 buyer behavior research found that AI chatbots are now the #1 source influencing vendor shortlists, and 79% of software buyers say AI search has changed how they research. TrustRadius’s 2025 report adds texture to the frequency: among respondents using AI at work, 95% use it at least weekly and 69% use it daily, and early research is cited as the most common AI use case in buying. That’s not a trend to track—that’s the current state of play.

What’s also shifting is where your content shows up in that research process, and whether it shows up at all. Pew Research’s behavioral analysis found that when an AI summary appears in search results, users click traditional results at roughly half the rate they otherwise would: 8% vs. 15%. And clicks directly on sources cited inside the AI summary? About 1%. The interface layer between buyer questions and vendor websites is no longer just Google’s ranking algorithm — it’s an AI mediation layer, and it operates by different rules.

This report lays out three strategic shifts that define content marketing for B2B software and services in 2026, explains why original research has become a more important investment than ever in an AI-mediated buying world, and closes with ten specific actions—each tied to a 2025–2026 primary source — that you can move on now (yes I included the links in a list at the end, for those of us who remember AP college essay writing).

## Definitions and the New AI-Mediated Buying Reality

Before getting into strategy, it’s worth aligning on terminology—not because the acronyms matter for their own sake, but because AEO, AIO, and GEO are showing up in every content

team's planning conversations right now, often interchangeably and sometimes incorrectly.

*(Please note: if you're looking for an SEO definition here, I'll refer you to the majority of software solutions in the space such as SEMrush, Ahrefs, Moz, and others for that)*

**AIO (AI Overviews / AI Mode optimization)** is the emerging umbrella term for optimizing content so it can be selected and shown as a cited or supporting link inside Google's AI Overviews and AI Mode experiences. Google's own guidance is clear that there are no special requirements beyond strong fundamentals—but it also explains that AI features may use query fan-out and can surface a "wider and more diverse set of helpful links" than classic search. In practical terms, your content must be easy to parse, trust, and cite.

**AEO (Answer Engine Optimization)** focuses on becoming the best answer for AI systems that provide direct responses—AI assistants, AI-infused search, and answer-style formats. The academic literature frames this as an emergent optimization field that contrasts with traditional SEO, because generative answers can draw from different source ecosystems than Google's ranked links.

**GEO (Generative Engine Optimization)** is the research-backed discipline of improving visibility within generative AI response systems that produce synthesized, citation-backed answers instead of a ranked list. A 2025 arXiv paper explicitly defines this shift and argues it necessitates a "new paradigm" beyond established SEO, because generative engines differ in sourcing behavior, freshness sensitivity, and sensitivity to phrasing.

The buying context that makes all three acronyms boardroom-relevant is no longer speculative. Gartner's 2026 data shows 67% of B2B buyers prefer rep-free experiences and 45% actively used AI in a recent purchase. G2's 2025 buyer research confirms AI chatbots have become the #1 source influencing vendor shortlists, and 79% say AI search changed how they research software. TrustRadius's 2025 report fills in the operational picture: 95% of AI-at-work users rely on it at least weekly, and early research is the most commonly cited use case in buying. Gartner explicitly warns sellers that static collateral won't cut it anymore and recommends structuring content into modular building blocks that can be dynamically assembled for workflow-fit support. That's a content operations problem as much as it is a strategy problem.

## **Three Strategic Shifts Content Marketers Must Make for AI-Driven B2B Buyers**

These aren't aspirational shifts for 2027 planning. For B2B software and services content teams, they're the minimum viable transformation for staying relevant as buyers route more of their evaluation through AI.

---

### **Shift from Traffic-First to Answer-First Visibility**

The classic B2B content model was built on a sequential funnel: publish → rank → click → convert. The logic was sound for its time. In 2026, buyers frequently get an answer—or a shortlist—before they ever visit your site. And even when they do encounter your content, it may reach them filtered through an AI summary they never clicked through to actually read.

Pew's behavioral analysis makes the stakes concrete: when AI summaries appear, click rates on traditional results drop to **8% from 15%**, and clicks on sources cited inside the AI summary hover around **1%**. That's not a rounding error—it's a fundamental shift in how attention flows from a search result page. Getting "included" in an AI answer is becoming a visibility surface entirely distinct from earning a blue-link click, and it requires a different optimization strategy.

Google's documentation explains the mechanism: AI Overviews and AI Mode can expand retrieval via query fan-out to generate a response and discover supporting pages, then surface links as supporting evidence and exploration paths. Conductor frames this as a "parallel surface of visibility" where brands are seen within AI answers before any click occurs.

The practical implication: keep measuring traffic, but upgrade your reporting to include AI inclusion signals—citations, mentions, share-of-voice in AI answers—plus assisted pipeline impact. The first impression may be happening off-site, without a click, and most current attribution models aren't capturing it.

---

## Shift from Commodity Publishing to Primary-Source Authority

Generative AI systems are extremely good at synthesizing generic explanations. That makes undifferentiated, keyword-optimized content a depreciating asset—not because the content is inherently bad, but because AI can produce an adequate version of it in seconds and buyers increasingly know it. Google's people-first guidance is unusually direct here: ***it asks whether content provides "original information, reporting, research, or analysis,"*** and whether it adds ***substantial originality instead of rewriting others*** (this last part is important for later). In AI-mediated discovery, originality isn't a virtue—it's an eligibility strategy.

Practitioner guidance is converging on the same point. Backlinko's 2025 AI Overviews guide lists "publish original data" as a core optimization move and asserts that AI systems favor stats, benchmarks, and transparent pricing. For B2B software and services, that translates into research-led assets—benchmarks, implementation learnings, ROI calculators with stated assumptions, security and compliance explanations with actual evidence—that buyers can reuse in internal justification cycles.

What's actually happening is that content needs to behave more like market intelligence: defensible, specific, sourced, and designed to be cited. AI layers surface what appears most reliable and reusable, not what is most promotional. The only way to build a durable position in this environment is to ***own a piece of evidence that others have to come back to you to cite.***

---

## Shift from Human-Only Storytelling to Machine-Readable Content Operations

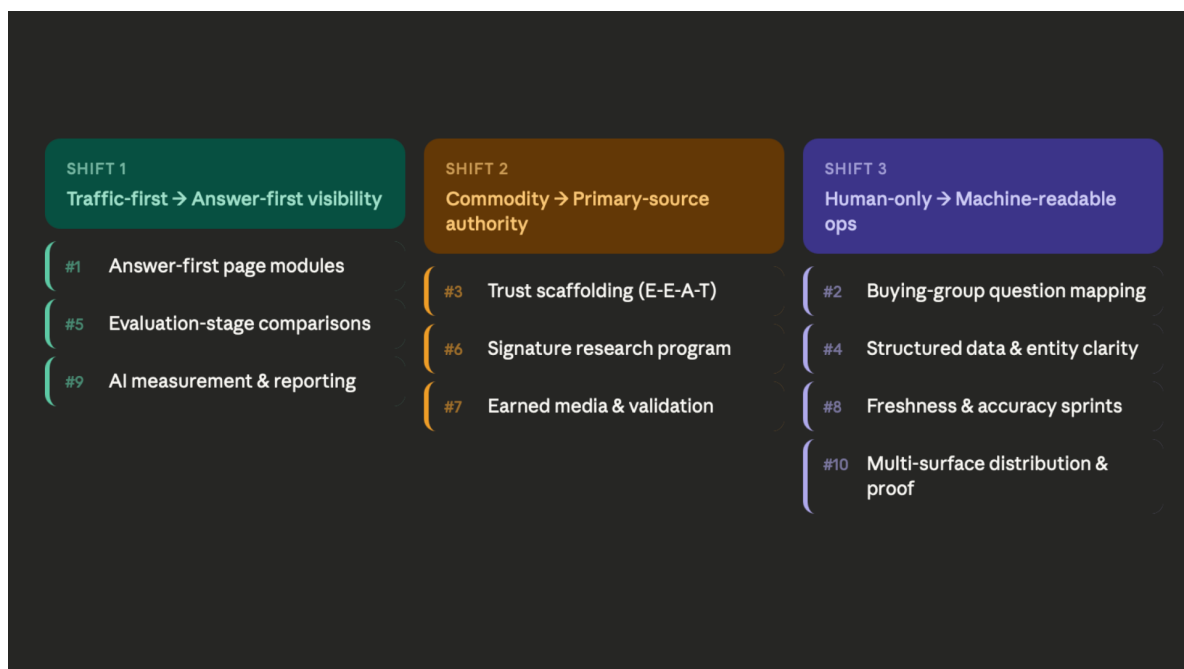
This is the shift most content teams are slowest to act on, because it feels like a technical problem rather than a content problem. It's both.

Microsoft's guidance for AI search selection describes how AI assistants "parse" content into smaller structured pieces evaluated for authority and relevance, then assemble an answer from multiple sources. That means your information architecture—headings, modular sections,

tables, Q&A blocks, consistent terminology—is now a direct input to AI selection systems, not just a formatting preference for human readers.

Google's AI features guidance gives the practical checklist: allow crawling, make content findable through internal links, ensure important content is available in textual form, and ensure structured data matches visible text. It's not exciting guidance, but the implication is serious: if your content is fragmented, duplicative, or structured around a 2018 editorial workflow, AI systems will produce inconsistent summaries of who you are and what you do. Brand narrative drift at scale is what happens when content operations aren't treated as infrastructure.

The diagram below maps each of these three strategic shifts to the specific tactical actions covered later in this report.



## What "Good" Looked Like Before AIO/AEO/GEO Changed the Interface

Pre-AI, good content marketing was largely a disciplined publishing and demand-capture machine. The Content Marketing Institute's definition—creating and distributing valuable, relevant, consistent content to attract and retain a defined audience and drive profitable customer action—was and remains strategically sound. What's changed is the operating assumptions sitting underneath it.

Google's SEO Starter Guide describes SEO as helping search engines understand content and helping users find your site and decide whether to visit. That interface model implicitly optimized for clicks as the primary unit of value. Featured snippets foreshadowed today's answer-engine world—Google could "reverse the format" of a result and show the snippet first, effectively answering the query on the SERP—but snippets still lived inside a click-forward results page, and winning the snippet often increased traffic and authority simultaneously.

In short, pre-AI good content marketing was optimized for discoverability via rankings, persuasion on owned pages, and conversion pathways that assumed the buyer would land on your site early enough to be influenced there. That model isn't obsolete—but it's no longer sufficient on its own.

Here's how the two eras compare across the dimensions that matter most for B2B content strategy:

| Dimension                 | Pre-AI / pre-AEO                              | AI era (AIO/AEO/GEO)   |
|---------------------------|---|--|
| Primary visibility target | Ranked listings → clicks to owned site        | Inclusion in AI answers + supporting links; diversified link set via query fan-out |
| Core success metric       | Sessions, CTR, form fills, MQLs               | Citations/mentions/share-of-voice in AI answers, assisted conversions              |
| Content win condition     | "Best page on the topic" + strong on-page SEO | "Best extractable fragment" + trust signals + machine-readability                  |
| Buyer attention pattern   | Higher likelihood of clicking to sources      | Lower click propensity when AI summary appears; ~1% click on AI-summary sources    |
| Strategic content moat    | Topic clusters; backlink acquisition          | Primary-source data + third-party validation + structured, modular content ops     |

## Why Original Research Matters More Now Than It Ever Did

Andy Crestodina at Orbit Media has been making the argument for brands to establish themselves as the source of original, 1st party data-backed research for years, and in 2026, it's never been more relevant. Original research tends to "crush" other content types because it makes you the primary source—earning links, shares, and credibility that no amount of well-optimized explainer content can match. His framing in the Orbit Media article on original research: research wins because it is new and unique and because it makes your site the source others cite. That was true before AI-mediated search. It's significantly more true now.

What changed in 2025–2026 is that original research doesn't just outperform in traditional distribution channels. It becomes a defensive asset in AI-mediated discovery—your "receipts" that survive summarization, procurement scrutiny, and the buyer's internal AI workflows. Three recent sources sharpen the ROI case:

1. Google's 2025 content quality guidance explicitly rewards originality by asking whether content provides "original information, reporting, research, or analysis," and whether it avoids simply rewriting other sources. If you want Google—and by extension, Google's AI features—to treat your pages as high-value inputs, being the originator rather than the remix is exactly what the evaluation rubric rewards.

2. Microsoft's 2025 guidance for AI search selection takes this further in operational terms: it recommends avoiding vague language and anchoring claims in measurable facts. In content marketing language, that's a mandate for evidence-led narratives: benchmarks, quantified outcomes, explicit assumptions, verifiable data. Original research is the most scalable way to generate measurable facts that belong to you and only you.

3. Backlinko's 2025 AI Overviews playbook makes the recommendation plain: "Publish original data," noting that AI systems favor stats, benchmarks, and transparent pricing. That's the content moat generative systems can't conjure from consensus alone—and it's the kind of asset other sites reference, reinforcing your authority across both AI and traditional search surfaces.

The through-line: in a world where the interface summarizes everything, original research is how you keep your brand from becoming a footnote in someone else's answer.

## Ten Content-Marketing Moves to Improve AIO/AEO/GEO Now

- **1. Reformat priority pages into answer-first modules:**
  - Design every high-value page so the first "screen" answers the core question in plain language, then expands with supporting sections that are individually extractable—Q&A blocks, short definitions, tables, bulletproof examples. Microsoft's guidance explains that AI systems parse pages into smaller pieces and that direct Q&A pairs, clear headings, lists, and tables are formats assistants can lift directly into answers. Modularity isn't just UX polish—it's selection readiness. (Reference: Microsoft Ads Blog—"Optimizing Your Content for Inclusion in AI Search Answers," October 2025.)
- **2. Map the editorial roadmap to buying-group questions, not just keywords:**
  - Modern B2B deals are decided by committees doing rep-light research, often with AI assist. Build content "lanes" aligned to stakeholder tasks: IT/security validation, finance value clarity, implementation and change management, procurement comparison, and risk mitigation. The 6sense 2025 Buyer Experience Report states that buyers need to understand how AI is embedded in solutions—capabilities, pricing, implementation timelines, and data security—and that nearly all buyers reported using LLMs to research solutions. Your content needs to answer cross-functional due diligence questions, not just top-of-funnel curiosity. (Reference: 6sense—2025 Buyer Experience Report.)
- **3. Turn E-E-A-T into visible trust scaffolding—not a vibes-based aspiration:**
  - Add explicit sourcing, author expertise context, editorial review notes (especially on YMYL-adjacent risk topics like security), and clear "what changed" versioning on cornerstone pages. Google's guidance highlights that evaluators should ask whether content inspires trust through clear sourcing and evidence of expertise, and explicitly asks whether content provides original research and analysis rather than rewrites. This is how you reduce AI-era brand narrative drift: make trust machine-visible and buyer-visible at the same time. (Reference: Google Search Central—"Creating helpful, reliable, people-first content.")
- **4. Standardize structured data and entity clarity across your core content library:**
  - Treat schema as content packaging: consistent Organization/Product/SoftwareApplication/FAQ markup clarifies what a page is and reduces ambiguity in how systems interpret it. Google's structured data guidelines (updated 2026) emphasize that structured data must be representative of the page's main content, not misleading, and must match what's visible to readers. Schema strategy should be governed like editorial quality—not bolted on as a technical afterthought. (Reference: Google Search Central—"General structured data guidelines.")
- **5. Build "AI comparison-ready" assets for evaluation-stage prompts:**

- Assume buyers will ask AI to compare vendors, categories, and approaches — "X vs Y," "best for our constraints," "alternatives to Z." Publish comparison pages and decision guides that are factual, regularly updated, and structured for scanning. G2's 2025 buyer research reports that AI chatbots influence vendor shortlists and that 79% of software buyers say AI search changed how they research; it also explicitly advises sellers to promote competitive comparisons and proof points to make shortlists. That's your content brief. (Reference: G2—2025 Buyer Behavior Report.)
- **6. Launch a signature research program designed to produce reusable "proof blocks":**
  - Pick one research motion you can repeat: annual benchmark, quarterly pulse, implementation dataset, ROI model with transparent assumptions. Publish it in HTML as the canonical source, then atomize it into charts, snippets, and "stats with context" modules distributed across the site. Backlinko's 2025 AI Overviews guidance explicitly calls out publishing original data—stats, benchmarks, transparent pricing—as a lever for AI-era visibility, because it's the easiest thing for AI systems and buyers to reuse and cite. (Reference: Backlinko—"AI Overviews: The Definitive Guide," 2025.)
- **7. Make earned media a first-class distribution channel for AI credibility:**
  - In AI sourcing ecosystems, third-party validation can matter more than self-assertion. Treat digital PR, partner publications, analyst-style contributed pieces, influencer campaigns, and credible review ecosystems as part of the content engine—not a separate team's nice-to-have. Academic evidence from 2026 shows AI answer engines can have very low domain overlap with Google's top results, and that generative systems systematically privilege certain source types—often earned and authoritative—in ways that differ meaningfully from traditional search. You win by ensuring your best evidence exists in places AI systems already trust. (Reference: arXiv preprint, 2026—on AI answer engine sourcing behavior.)
- **8. Run a quarterly freshness and accuracy sprint on "AI-likely" pages:**
  - Prioritize pages that answer category-defining questions, core pain points, and high-stakes risk topics—these are the pages AI tools most often pull into summaries. Update claims, screenshots, and "last verified" notes; retire stale pages that create contradictions. Google's structured data guidelines explicitly call out providing up-to-date information and note that Google may not show rich results for time-sensitive content that is no longer relevant. Freshness isn't merely a ranking factor—it's a selection filter for systems trying to avoid citing outdated information. (Reference: Google Search Central—"General structured data guidelines.")
- **9. Upgrade measurement from "rankings + leads" to AI visibility and pipeline influence:**
  - Instrument reporting that tracks where your domain is cited or mentioned inside AI experiences, how AI referral traffic behaves, and what content gets pulled into AI Overviews. Conductor's 2026 benchmarks operationalize this by quantifying AI referral traffic, AI citation and mention market share, and the share of Google searches triggering AI Overviews—framing AI visibility as a distinct performance channel where credibility, not clicks, is the leading indicator. (Reference: Conductor—"AEO/GEO Benchmarks Report," 2026.)
- **10. Treat reviews, UGC, and proof surfaces as content—then integrate them into your narrative:**
  - In AI-mediated buying, buyers fact-check and triangulate. Build a content program that incorporates customer voice: review response content, case evidence, community Q&A, and "what customers say" sections embedded in your key pages — not buried in a resources folder. TrustRadius's 2025 report explicitly states that early research is where AI is taking over from traditional search, includes vendor guidance to prioritize buyer-question content to increase chances of being cited by an LLM, and emphasizes that trusted content—including UGC—helps buyers fact-check AI tools. (Reference: TrustRadius—"From Buzzword to Backbone," 2025.)

## Sources

1. Google Search Central—AI Features and Your Website: <https://developers.google.com/search/docs/appearance/ai-features>
2. Pew Research—Google users click less when AI summary appears (July 22, 2025): <https://www.pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results/>
3. Google Search Central—Creating Helpful, Reliable, People-First Content: <https://developers.google.com/search/docs/fundamentals/creating-helpful-content>
4. Microsoft Ads Blog—Optimizing Your Content for Inclusion in AI Search Answers (Oct 2025): <https://about.ads.microsoft.com/en/blog/post/october-2025/optimizing-your-content-for-inclusion-in-ai-search-answers>
5. arXiv—Academic preprint on AEO/GEO (2601.16858): <https://arxiv.org/pdf/2601.16858>
6. arXiv—2025 GEO new paradigm paper (2509.08919): <https://arxiv.org/abs/2509.08919>
7. Gartner—2026 Sales Survey: 67% of B2B Buyers Prefer Rep-Free Experience: <https://www.gartner.com/en/newsroom/press-releases/2026-03-09-gartner-sales-survey-finds-67-percent-of-b2b-buyers-prefer-a-rep-free-experience>
8. G2—2025 Buyer Behavior Report: <https://research.g2.com/cm05-2025-buyer-behavior-report-research-g2>
9. TrustRadius—From Buzzword to Backbone (2025): <https://go.trustradius.com/rs/827-FOI-687/images/From%20Buzzword%20to%20Backbone.pdf>
10. Conductor—AEO/GEO Benchmarks Report (2026): <https://www.conductor.com/academy/aeo-geo-benchmarks-report/>
11. Backlinko—AI Overviews: The Definitive Guide (2025): <https://backlinko.com/ai-overviews>
12. Content Marketing Institute—What Is Content Marketing: <https://contentmarketinginstitute.com/what-is-content-marketing>
13. Google Search Central—SEO Starter Guide: <https://developers.google.com/search/docs/fundamentals/seo-starter-guide>
14. Google Search Central—Featured Snippets: <https://developers.google.com/search/docs/appearance/featured-snippets>
15. Orbit Media—5 Examples of Original Research in Content Marketing: <https://www.orbitmedia.com/blog/research-best-form-of-content/>
16. 6sense—2025 Buyer Experience Report: <https://6sense.com/science-of-b2b/buyer-experience-report-2025/>
17. Google Search Central—General Structured Data Guidelines: <https://developers.google.com/search/docs/appearance/structured-data/sd-policies>