

CRITERION · 2026 · NO. 01

Cited or Invisible

The AI Citation Intelligence Report

How AI engines decide which B2B vendors exist, and which ones don't.

AI ENGINES DON'T RANK. THEY CITE.

THE PROBLEM

The buying decision moved. Most brands didn't notice.

B2B buyers used to Google it. Now they ask ChatGPT.

The question is the same. The answer is not.

AI systems don't return a list of links. They return a verdict. One answer. A short list of vendors. Maybe a citation or two.

If a brand isn't in that answer, it doesn't exist for that buyer. Not page two. Not below the fold. Gone.

This research set out to answer one question: who is actually getting cited inside AI buying answers, and why.

NOTE

The patterns identified in this research led to the development of Criterion, an AI citation intelligence platform currently in beta.

METHODOLOGY

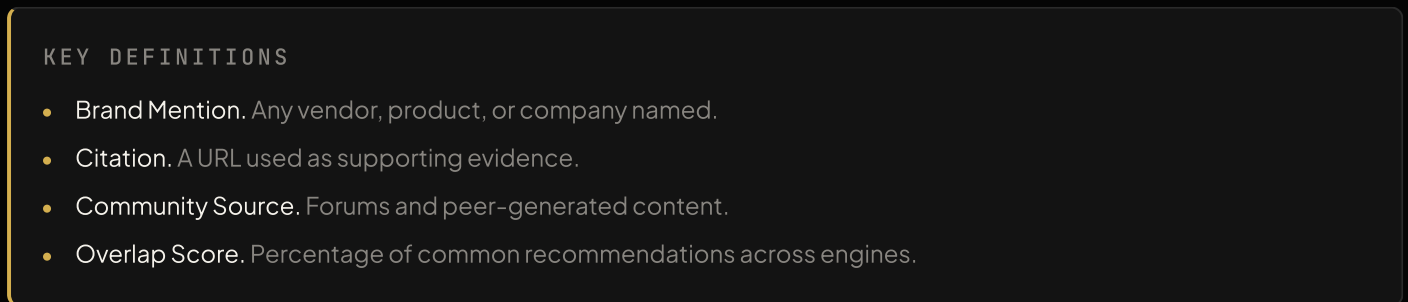
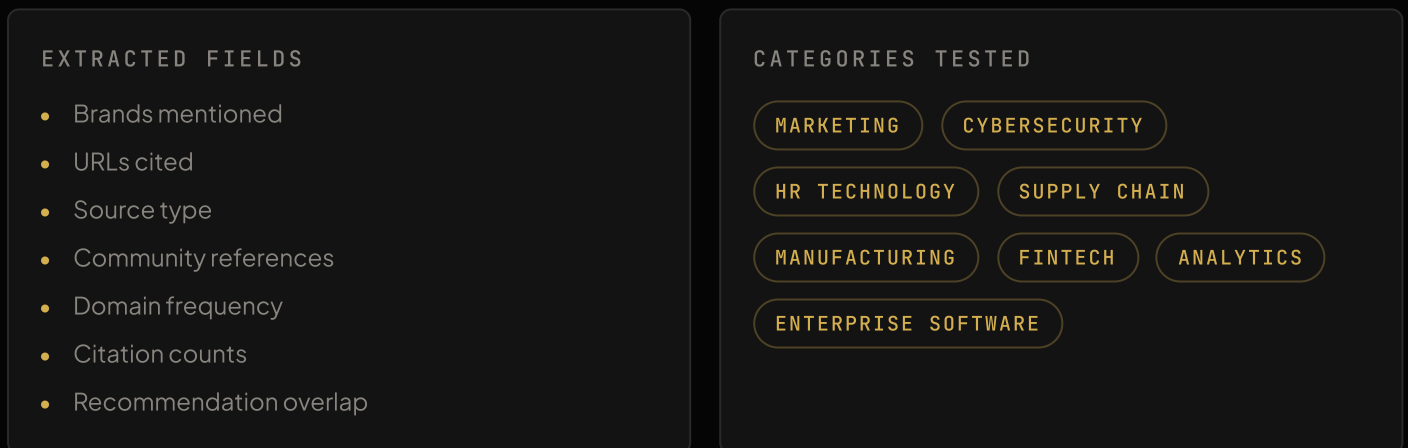
How this study was built.



Every prompt represented evaluation-stage buying behavior. The kind a buyer sends when ready to shortlist vendors, not just explore.

110 prompts were tested across ChatGPT, Claude, and Gemini. Each answer was analyzed for every brand mentioned, every URL cited, every domain referenced.

The result: 329 answers, 3,266 brand mentions, 1,375 cited URLs across 697 unique domains, and a normalized dataset exceeding 6,000 records.



FINDING 01

Being named is not the same as **being trusted**.

60.6%

MENTIONED BUT NOT CITED

60.6% of brand mentions had no URL behind them. AI recalled the name. It did not cite the proof.

For smaller brands, citations were the only path in. Familiarity without evidence is a dead end.

IMPLICATION

The brands winning AI aren't the most famous. They're the most cited.

Large vendors often benefit from familiarity. Smaller brands appear more often only when supported by sources. Recall and trust travel through different mechanisms.

Mentioned

Named in the answer.

Cited

Backed by a URL.

Trusted

Cited across multiple engines.

FINDING 02


Reddit beat Gartner. Nobody planned for this.

Reddit was the most cited domain in the dataset. Not G2. Not Gartner. Not vendor blogs. Reddit.

LinkedIn, the platform every B2B marketer lives on, represented 0.9% of all cited URLs.

The implication is uncomfortable. Peer threads are shaping AI recommendations more than analyst reports. Community is not a nice-to-have. It is the citation layer.

TOP-CITED DOMAINS · SHARE OF DATASET



Counts represent leading-domain frequency in the cited URL set. LinkedIn appears as a percentage of all citations to highlight its absence relative to its B2B reputation.

IMPLICATION

LinkedIn thought leadership is not reaching AI.

A Reddit presence might be. The standard B2B playbook is missing the surface that actually feeds the citation layer.

FINDING 03

There is no single AI citation graph. There are three.

The same buying question. Three different engines. Three almost entirely different sets of vendors and sources.

Only 3.0% of cited domains appeared across all three engines. 15.6% of questions had zero brands in common.

Optimizing for one engine and calling it AEO is like doing SEO for one search engine and ignoring the others.

ChatGPT

5.7%

OVERLAP WITH CLAUDE

- Broader synthesis
- Structured answers
- Moderate source counts

Claude

2.1%

OVERLAP WITH GEMINI

- Larger source sets
- Nuanced explanation
- Higher citation density

Gemini

7.6%

OVERLAP WITH CHATGPT

- Concise retrieval
- Fewer citations
- Higher uncited-brand rates

3.0%

SHARED ACROSS ALL THREE ENGINES

A brand cannot optimize for a single AI engine.

Each engine pulls from a largely disjoint citation set. Buyers see a different vendor list depending on where they ask.

FINDING 04

Comparison pages are the highest-leverage asset in AI search.

Comparison and alternatives pages made up 29.2% of all citations across the dataset. In comparison-style prompts, that number jumped to 56.3%.

AI systems are not surfacing the about page. They are surfacing the page that answers: why this product over that one.

CONTENT PATTERN	BEHAVIOR	CITATION SHARE
Comparison, versus, alternatives	Structured decision frameworks	29.2%
Comparison prompts → comparison pages	Format-matched retrieval	56.3%
Vendor documentation	Primary-source authority	High
Thought-leadership posts	Narrative, low structure	Low

IMPLICATION

Most brands do not have one.

Comparison-style content remains rare in B2B catalogs, even though it dominates the citation layer for evaluation-stage prompts. The gap is the opportunity.

FINDING 05

Your docs are doing more work than **your blog.**

Technical documentation consistently outperformed thought leadership in AI citation rates.

In cybersecurity and manufacturing, vendor docs and standards bodies dominated the citation layer. In marketing and supply chain, review platforms and analysts carried more weight. Category matters. Content strategy needs to match it.

CYBERSECURITY

Standards-heavy

Government and standards bodies dominate.

MANUFACTURING

Docs-heavy

Vendor documentation carries the citations.

MARKETING · SUPPLY CHAIN

Review-heavy

Review platforms and analysts lead.

IMPLICATION

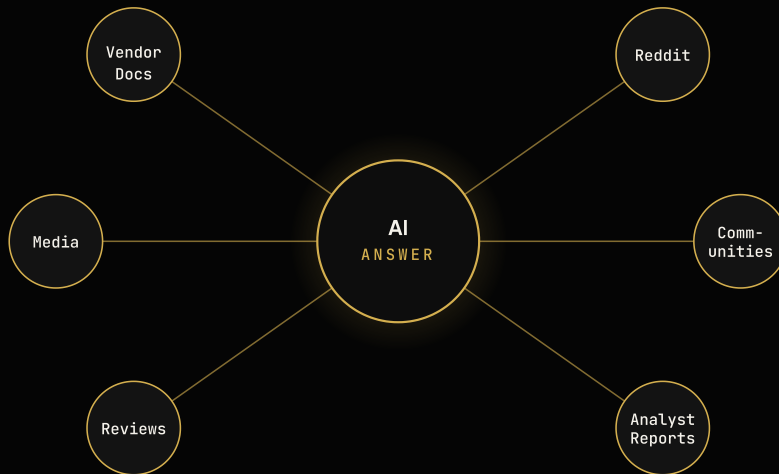
Citation channels are category-shaped.

A playbook that works in marketing rarely transfers cleanly to cybersecurity. Brands that match content type to category citation behavior consistently outperform those that don't.

FRAMEWORK

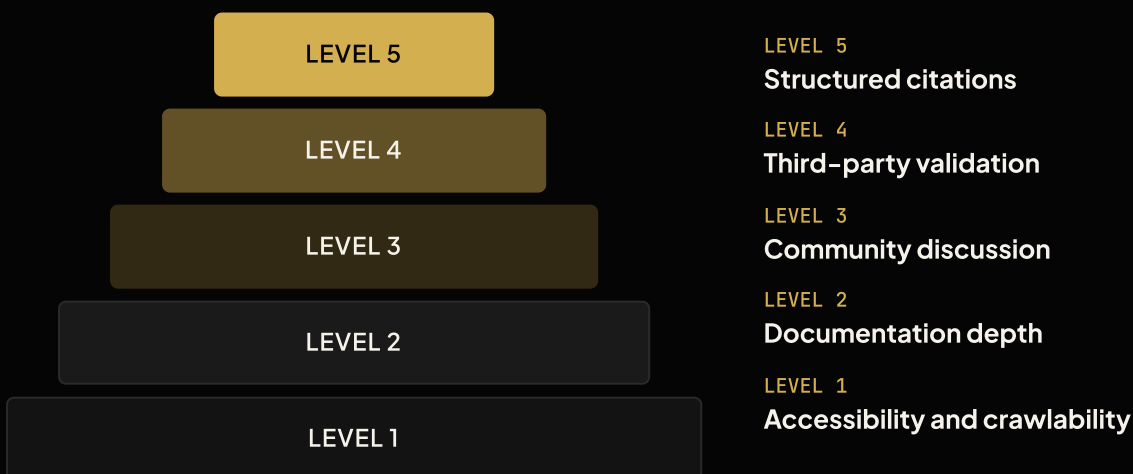
AI doesn't have opinions. It has sources.

Every AI answer is a synthesis of whatever it could find and trust. No single channel dominates. The brands that appear across the most layers win the most answers.



AI trust is built in layers.

Brands that occupy multiple layers get cited consistently. Brands that occupy only one or two get recalled without evidence, or not recalled at all.



BETA VALIDATION

The research predicted it. A real company proved it.

During the research period, a B2B company in the space management technology category participated in structured beta testing. The findings from this study were applied directly. Here is what happened.

WHAT THE AUDIT FOUND

- AI was citing wrong pricing from Reddit threads.
- Brand information was inconsistent across sources.
- The company had almost nothing AI could cite.
- No content existed in citation-ready format.

WHAT WAS DONE

- Incorrect Reddit pricing threads were addressed.
- New educational content was published for AI retrieval.
- Citation-friendly documentation was expanded.
- Visibility gaps identified in the audit were filled.

VISIBILITY PROGRESSION



72.4%

INCREASE IN AI VISIBILITY OVER 90 DAYS

"Smaller B2B brands are the ones most hurt by AI invisibility — and the least able to pay enterprise rates to fix it."

PRACTICAL IMPLICATIONS

What to do **this week.**

FOR MARKETERS

Build pages AI can cite.

- Write comparison pages for every major competitor.
- Build alternatives pages for the category.
- Create implementation guides with real specificity.
- Write FAQs that answer exact buyer questions.

FOR SEO TEAMS

Stop tracking only rankings.

- Track which domains AI is citing in the category.
- Monitor citation frequency, not just position.
- Measure community visibility on Reddit and forums.
- Watch how different engines treat the same brand.

FOR PRODUCT MKTG

Docs are a distribution channel.

- Write how-to content built for AI retrieval.
- Create use-case decision frameworks buyers quote.
- Tag claims with schema so AI can parse them.
- Keep pricing and capability pages current.

A common thread runs across all three roles. Citation visibility is built by structured, machine-readable content that answers specific buyer questions. The teams that adapt fastest treat AI retrieval as a first-class distribution surface.

PREDICTIONS

What happens **by 2027.**

BY 2027 • CHANNEL

AI assistants become the first stop for B2B vendor research. Not a supplement to search. The starting point.

BY 2027 • INFRASTRUCTURE

Citation visibility becomes a tracked, managed surface. The same way SEO became a discipline, AEO becomes a function.

BY 2027 • COMMUNITY

Peer content carries more weight in AI answers. Community presence stops being optional for B2B brands.

BY 2027 • RANKINGS

Search rank stops predicting AI visibility. The two become separate problems requiring separate strategies.

None of these shifts arrive at once. The transition mirrors the early SEO years: a slow build, then a sudden inflection. The brands that begin tracking citation behavior now will have a measurable head start when the inflection arrives.

CONCLUSION

The game changed. The scoreboard changed. Most brands are still playing the old game.

This research analyzed 329 AI-generated B2B buying answers and found a consistent pattern.

AI systems do not treat all brands equally. They treat cited brands differently from recalled ones. They treat community-validated brands differently from undocumented ones. They build almost entirely separate citation sets depending on which engine a buyer uses.

60.6% of brand mentions had no citation behind them. Reddit outranked every analyst platform in citation frequency. Only 3% of domains appeared across all three engines.

The brands that understand this will build for citation, not just visibility. The ones that don't will keep optimizing for a signal that no longer predicts where buyers find them.

SHIFT 01

From positions to citations.

The unit of B2B visibility is no longer a rank. It is a URL inside an AI answer.

SHIFT 02

From channels to sources.

No single channel controls AI visibility. Every layer of evidence compounds.

SHIFT 03

From reporting to action.

Knowing the problem is not enough. The work is fixing what is pulling visibility down.

AI engines don't rank. They cite.