

# The Hidden Selection Phase

How B2B Brands Win or Lose in AI  
Before Sellers Ever Engage



A Foundation × AirOps Research Report

50 Brands | 7 Verticals | 5.1 Million AI Responses | 57.2 Million Citations

# Executive Summary

Before a single sales call is booked, before a demo is scheduled, or a buyer even visits your website, there is a moment that increasingly determines whether your brand makes the shortlist or gets left off entirely. We call it the **Selection Phase**. And today, it runs through large language models.

This report examines how 50 B2B brands across seven verticals appear, get cited, and get described inside AI responses. Drawing on 5.1 million responses and 57.2 million citations tracked across five platforms over 60 days, Foundation and AirOps map the gap between how brands see themselves and how AI represents them to buyers.

**Finding: only 10.15% of all citations in LLM responses point to brand-owned domains.** The remaining roughly 90% come from third-party sources such as review sites, Reddit threads, YouTube videos, competitor pages, and community forums — that brands do not control. At the response level, 68% of AI answers cite only external sources. In other words, in more than two out of three responses, the brand's own content is completely absent.

This matters because the buyer's journey has compressed. Thirty percent of B2B buyers now start their research with AI before Google, and 94% of buying groups are using LLMs at some point in their journey. And by the time buyers do reach out to sales, 94% have already ranked their shortlist, with the first-choice vendor winning 80% of the time.

That shortlist has shrunk from four or five vendors to two or three. Missing the AI-mediated selection phase could exclude you entirely. We examined this gap through three dimensions:

## **The Attribution Gap.**

Who introduces your brand to AI users? Across all 50 brands, only 10.15% of citations came from brand-owned sources. For unbranded discovery queries, where buyers explore a category without naming a specific vendor, that number drops to 2.24%.

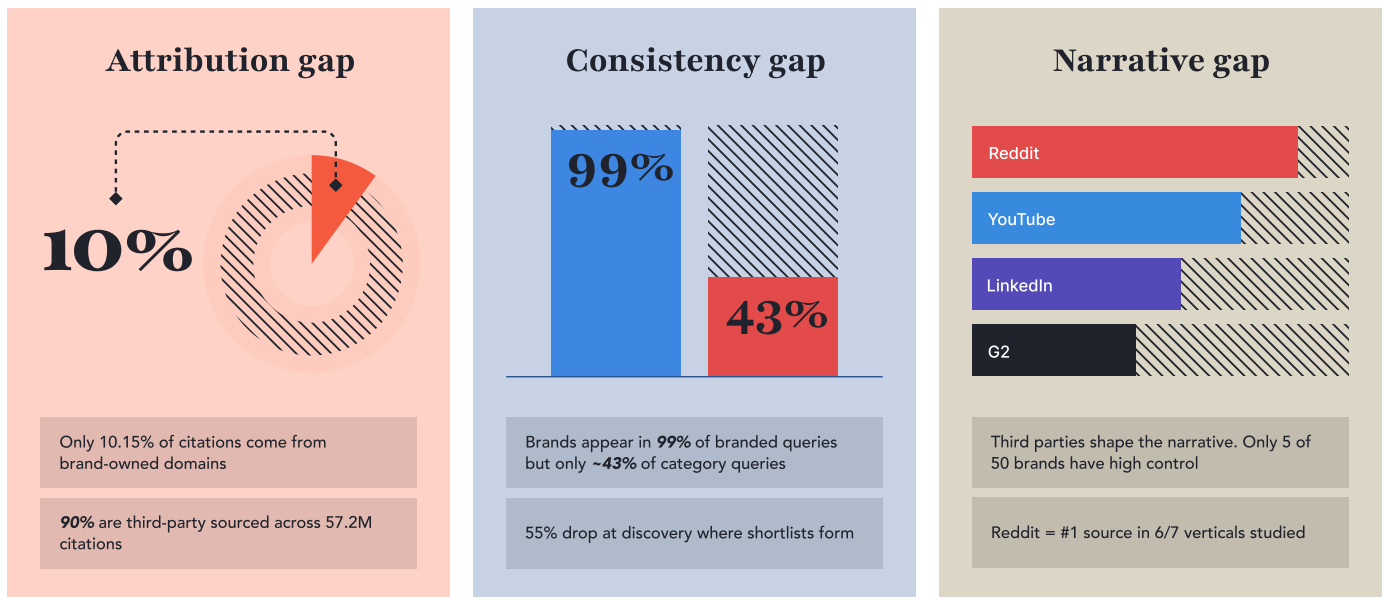
## **The Consistency Gap.**

How reliably does your brand appear? Brands show up in roughly 99% of AI responses when users ask about them by name. But in category-level queries, the ones actually forming shortlists, the average mention rate falls to 44.6%. That is a 55-percentage-point drop between being asked about and being discovered.

## **The Narrative Gap.**

When your brand does appear, which sources shape the story — and does their framing reflect yours? Reddit and YouTube together account for roughly 34% of all external citations. The rest comes from review sites, professional networks, competitor domains, and community forums. Only five of 50 brands studied achieved what we define as high narrative control: owned content present in the citation mix alongside those third-party sources.

# B2B Brands Must Address Three Gaps to Increase AI Visibility



Source: The Hidden Selection Phase | Foundation x AirOps | 50 brands, 7 verticals, 5.1M AI responses, Dec 2025 – Feb 2026

FOUNDATION x airOps

Brands that perform well across all three dimensions share clear patterns:

Comprehensive help documentation that AI models cite at scale, owned content that competes alongside (not instead of) third-party sources, and consistent visibility across platforms and query types.

The top-performing brands from our analysis are structurally positioned for inclusion in the AI-mediated selection phase. Brands below threshold are structurally excluded. What follows is a detailed analysis of each gap, the patterns that emerge across verticals, and a framework for what B2B brands should do next.

## SECTION 1:

# AI Now Decides Which B2B Brands Make the Buyer's Shortlist

The B2B buying process is undergoing a fundamental shift. For two decades, the B2B buying process followed a relatively predictable path. A buyer had a problem. They searched for it on Google, clicked through a list of blue links, read some blog posts, downloaded a whitepaper, maybe watched a webinar. They repeated this process over weeks or months, gradually narrowing from a broad category that could help solve their problem to a shortlist of vendors to a final decision.

Marketers understood this path and built infrastructure around it: SEO strategies, content funnels, lead magnets, nurture sequences. The system worked because the journey unfolded incrementally, and each increment was an opportunity to influence. But that system is now fracturing, and LLMs are the reason.

## The Buyer's Journey Has Been Compressed

The shift is best understood as a three-tier evolution in how buyers access information.

### The first tier is traditional search.

A buyer types a question into Google, gets a list of links, clicks through, reading, returning to results, and searching again. The process is incremental by design, moving from top-of-funnel queries to mid (and bottom-of) funnel queries, as the buyer gathers information from dozens of sources over multiple sessions.

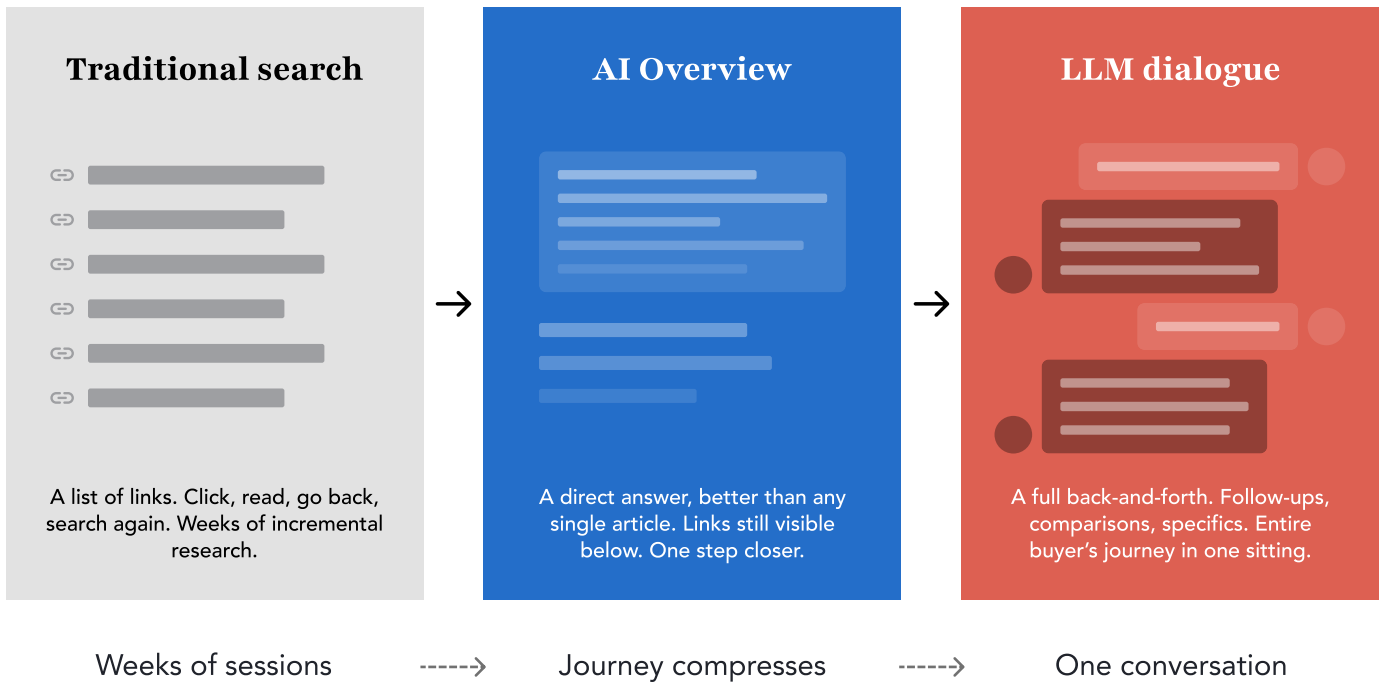
### The second tier is the AI Overview.

Google gives you a synthesized answer directly in the search results, often better than any single article. You still see the links, but the answer is already there.

### The third tier is the LLM dialogue.

A full back-and-forth conversation with an AI that can answer follow-up questions, clarify ambiguities, compare alternatives, and drill into the specifics of your situation. Instead of a single answer, the buyer gets a dialogue. And within that dialogue, they can complete in minutes the evaluation that used to require weeks of incremental searching.

## B2B Brands Must Address Three Gaps to Increase AI Visibility



A buyer who arrives at your website via an LLM referral has already done their homework. They've asked the AI which tools lead your category, how you compare to competitors, what your strengths and weaknesses are, and whether you are a good fit for their specific use case. What once took ten or fifteen search sessions can now happen in one conversation. By the time they click through to your site, they are approaching from the bottom of the funnel instead of the top.

# What This Means for B2B SaaS Brands

Buyer journey compression creates a cascade of consequences for B2B brands. Each one raises the stakes for the next.

## 1. Buyers start in AI before they start in Google.

According to a recent G2 report, 30% of B2B buyers now begin their research with AI before they ever open a search engine<sup>1</sup>. In enterprise organizations with 1,000+ employees, AI has already overtaken traditional search as the starting point for software evaluation. Seventy-nine percent of software buyers say AI search has changed their research process, and 89% report using generative AI across every phase of their buying process. Overall, 94% of buying groups are using LLMs. This means the first impression a buyer forms of your brand is increasingly shaped not by your website, your ads, or your sales team, but by what an AI model says when asked about your category.

## 2. Buyer shortlists are shorter than ever.

Buyers now evaluate two to three vendors, down from four or five. The consideration set has physically narrowed, and "no shortlist" behaviour — where buyers go directly to a single vendor — has doubled. Fewer slots means that missing the AI-mediated selection phase doesn't just reduce your chances of winning. It can exclude you from consideration entirely.

## 3. Those shortlists lock in before sales ever engages.

Ninety-four percent of buying groups rank their shortlist before ever contacting a seller. The first-ranked vendor wins 80% of the time<sup>2</sup>. The selection phase — the period when brands are evaluated, compared, and ranked — is largely over before sales teams know it has begun. If your brand didn't make the list during the AI-mediated research phase, there may not be a second chance.

**"AI is facilitating the buying process with or without your content, so you need to get involved."**


James Scherer, VP of Strategy, Foundation

# LLM Visitors Convert Better Than Google Visitors

The argument against AI search is volume. ChatGPT currently drives only a fraction of overall search traffic, and Perplexity even less. By raw referral numbers, LLMs are not yet competitive with Google. But volume is the wrong lens. Just look at the conversions Ahrefs drives through ChatGPT:

## Only a Fraction of Ahrefs Traffic Comes From ChatGPT, but It Drives 12% of Their Signups.

SHARE OF AHREFS TRAFFIC



of total website traffic comes from ChatGPT

24x

SHARE OF AHREFS SIGNUPS



of total signups come from ChatGPT

LLM referral traffic converts at 24 times higher efficiency. The volume is small. The signal is not.

Source: Ahrefs.com — "Does AI Search Traffic Convert Better Than Traditional Search?" by Patrick Stox, June 2025

FOUNDATION x aiOps

ChatGPT drives just 0.5% of Ahrefs' total website traffic<sup>3</sup>. It also drives 12.1% of their signups, roughly 24 times higher conversion efficiency than traditional search. The visitors arriving through LLM conversations are not casual browsers. They are buyers who have already evaluated, compared, and decided within the dialogue itself.

**"The referral traffic from LLMs has already gone through the buyer's journey. That's why it converts at dramatically higher rates — the searcher has done the entire funnel in a single dialogue with the LLM."**

James Scherer, VP of Strategy, Foundation

The traffic is small. The signal is not. If the buyer can complete their entire evaluation in a single LLM conversation, then the visitor who clicks through has already made their decision. They are not entering the funnel. They are confirming a choice.

# Why This Study Exists

**The Selection Phase now runs through LLMs.** Buyers start there, form shorter shortlists there, and lock in their preferences there, often before visiting a single vendor's website. The traffic that does come through converts at rates that dwarf traditional search.

And as company size increases, LLMs overtake Google as the primary source for software research entirely.

That creates a problem most B2B brands cannot currently solve: they have no visibility into how they are showing up in the AI-mediated selection phase. They do not know when they appear, what is being said, or where the AI is sourcing its information. The infrastructure that gave them influence over the traditional buyer's journey, SEO, content, funnels, does not map cleanly onto this new one.


To understand what brands are actually up against, we needed to measure it. This study sets out to do exactly that: to examine how B2B brands are appearing in LLM responses, where that sourcing comes from, and what separates the brands that make the shortlist from the ones that don't. That's the foundation for what follows.

## SECTION 2: WHAT WE DID

# Methodology

In partnership with AirOps, we measured how 50 B2B brands appear in AI responses and quantified the gap between brand-owned and third-party sourcing.

## Study Scope



60-day window December 22, 2025 - February 20, 2026

## Prompt structure 100 per brand 5,000 total



### Category + feature queries

- The prompt does not name a vendor
- Represents discover stage of AI journey
- Mirrors how buyers find potential solutions

### Comparison + evaluation queries

- The prompt mentions the brand by name
- Represents short list validation
- Mirrors how buyers assess a known vendor

*This split matters because branded and unbranded queries produce dramatically different results, as the data will show.*

## WHAT WE MEASURED

We classified every citation as brand-owned or external, and measured at two levels

### RESPONSE-LEVEL

**Did at least one citation point to the brand's own domain?**

Tells you how often your content shows up at all.

### CITATION-LEVEL

**Across all citations, what share points to the brand?**

Tells you how much of the overall sourcing your own.

---

These measurements feed three analytical dimensions

### Attribution gap

Who controls the sourcing?

### Consistency gap

How reliably do you appear across query types?

### Narrative gap

Who controls the story when you do appear?

*Brand names are anonymized throughout. Full methodology detail, including prompt structure, consistency scoring threshold, and sentiment classification, is available in the appendix*

# Only 10% of AI Citations Point to Brand-Owned Domains

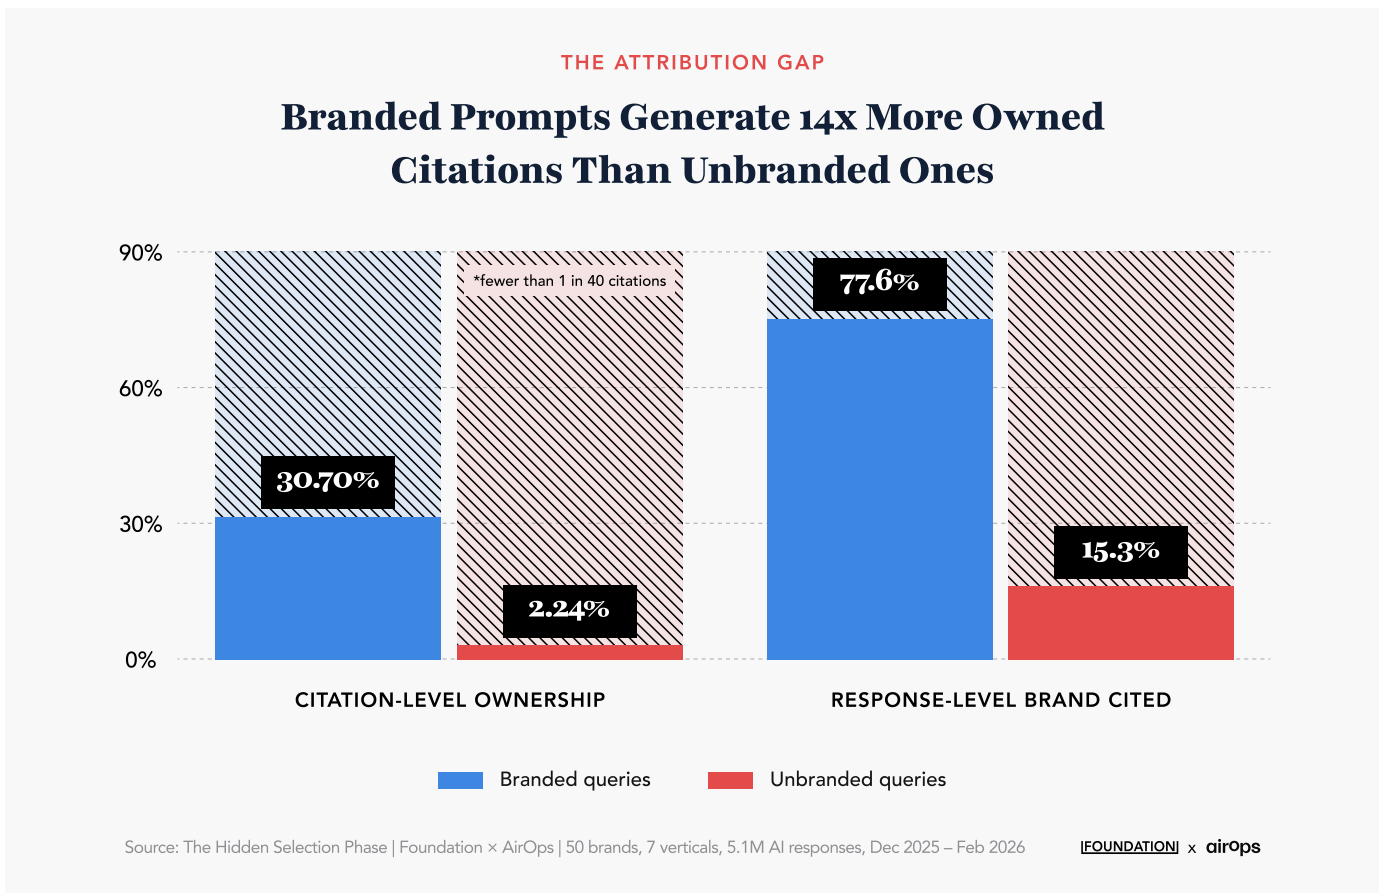
The Attribution Gap asks a simple question: when AI talks about your brand, where is it getting its information? Across 57.2 million citations, the answer is overwhelmingly clear. Not from your domain.

## 90% of AI Citations Come From Third-Party Sources

Across all 50 brands and seven verticals, only 10.15% of citations in LLM responses pointed to brand-owned domains. The remaining roughly 90% came from third-party sources: review sites, Reddit threads, YouTube videos, Wikipedia, competitor pages, news outlets, and community forums. At the response level, the picture is equally stark. Sixty-eight percent of all AI responses measured cite only external sources. This means that in more than two out of every three responses, the brand’s own content was completely absent from the sourcing — even when the brand itself was being discussed. Brands are being represented in AI, but they are not the ones doing the representing.

## When Buyers Search Without A Name, Only 2% of Citations Are Brand-Owned

The attribution gap widens dramatically depending on whether the buyer names a specific brand or explores a category.



**When a user asks about a brand by name, 77.6% of responses include at least one citation from the brand's own domain.** That sounds reassuring until you flip it and realize that in 22.4% of cases, even a branded query produces a response sourced entirely from third parties. A user asks about your product by name, and in approximately one out of five cases the AI answers using other people's content.

But the real gap appears in unbranded queries. When users ask about a category or feature without naming a brand, such as "What's the best marketing automation tool for mid-market?" or "Which DevOps platforms support Kubernetes natively?", only 2.2% of all citations point to brand-owned content.

**Only 15.3% of responses to unbranded prompts include any brand-owned citations.** This is the number that should concern every B2B marketer reading this report. Unbranded queries are discovery queries. They represent the moment a buyer is forming their initial consideration set, the moment that determines whether your brand makes the shortlist. And in that moment, 97.7% of the citations AI uses to shape the answer come from sources you do not own.

## The Extent of The Attribution Gap Varies by Vertical

The attribution gap exists in every vertical, but how big it is varies meaningfully. The figures below represent the aggregate across all query types, branded and unbranded combined.

Vertical	Brand-Owned %	External %
Workforce & Communications	15.18%	84.82%
Healthcare & Wellness	13.56%	86.44%
Fintech	11.43%	88.57%
Marketing & Analytics	10.40%	89.60%
DevOps & Security	8.46%	91.54%
Sales & Revenue	8.43%	91.57%
Productivity	6.30%	93.70%

Workforce & Communications leads with 15.18% brand-owned citations, which still means nearly 85% of sourcing comes from elsewhere. Productivity sits at the bottom with just 6.30%. The gap between the best-performing and worst-performing vertical is 2.4x, suggesting the problem varies in degree but is universal in presence.

**The takeaway:** The path to closing the attribution gap looks different depending on your vertical. What counts as citation-worthy content in Healthcare (medical institutions) is not the same as what counts in DevOps (developer documentation) or Productivity (integration guides). A vertical-specific content strategy is a must.

## Which Third Parties Dominate?

Across all 57.2 million citations, the external domain landscape breaks down into identifiable categories. Community sources, including Reddit, Medium, GitHub, Quora, and dev.to, account for approximately 27% of the top 50 external citation domains. YouTube represents roughly 13%, while LinkedIn accounts for about 11%. Competitor and adjacent SaaS domains make up approximately 15%, and brand help and support documentation accounts for about 8%, often hosted on subdomains that sit outside the primary domain.

Review sites like G2 and Gartner account for roughly 4%, less than many marketers might expect. Comparison and affiliate sites (TechRadar, NerdWallet, PCMag) account for about 3%. News and media (Forbes and similar) account for roughly 2%.

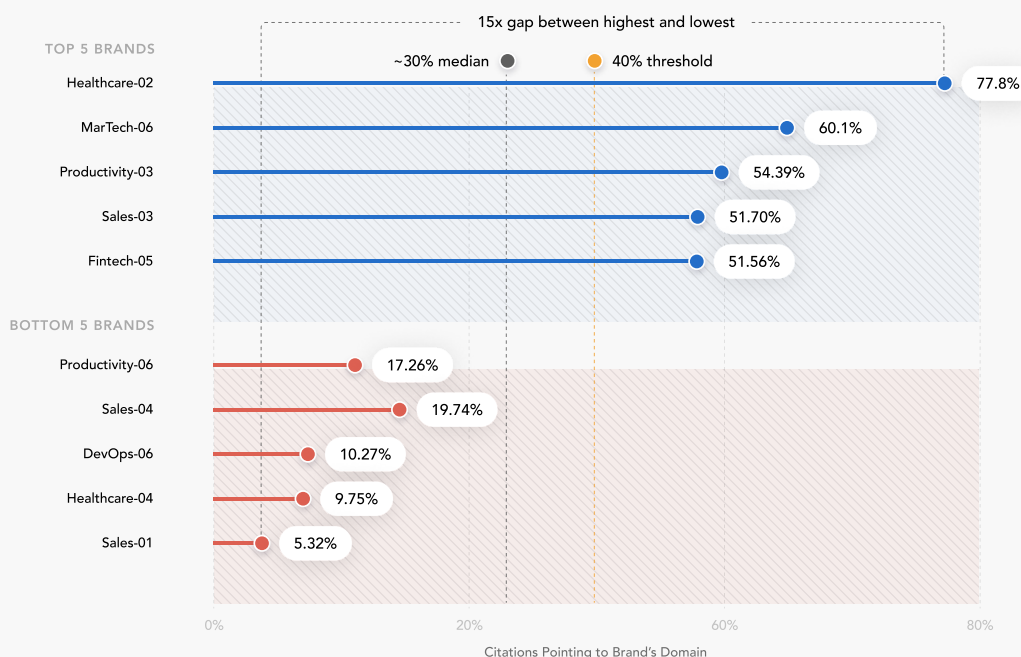
Community content and video dominate, while professional networks and competitor ecosystems also play a significant role. Traditional review sites remain important, but represent a smaller share than the overall community and content ecosystem.

## One Brand Owns 78% of Its AI Citations. Another Owns Just 5%.

The range of citation sources across the brands we analyzed is dramatic. The top-performing brand, Healthcare-02, had a response-level brand citation rate of 77.84%, meaning its own content appeared in more than three-quarters of all AI responses about it. The lowest, Sales-01, had just 5.32%. The median sits at approximately 30%.

### The Top-Performing Brand Receives 15x More Owned Citations than the Lowest

Brands above the 40% threshold have more consistency in MoM AI citations



Source: The Hidden Selection Phase | Foundation x AirOps | 50 brands, 7 verticals, 5.1M AI responses, Dec 2025 – Feb 2026 | FOUNDATION x airOps

This range matters because it proves the attribution gap is not outside of brands' control. Some brands have built content ecosystems that AI models cite heavily. By contrast, others in comparable verticals with comparable market positions have almost no owned presence in AI responses. The gap is wide, and it is movable.

Brands with response-level brand citation rates of 40% or higher had an average consistency score of 73.8, compared to 64.1 for those below the threshold. While this does not directly measure response inclusion, the correlation between citation ownership and consistency is clear: brands that own more of their citations also tend to appear more consistently present.

# Most Brands Only Show Up in AI When a Buyer Already Knows Their Name

The Attribution Gap tells you who controls your sources. The Consistency Gap measures how reliably your brand appears across the full range of queries buyers actually use. A brand might show up every time someone asks about it by name, but disappear when buyers explore the category or ask about specific features. That unevenness is the Consistency Gap, and it determines whether your brand makes the AI-mediated shortlist or only appears when someone already knows to ask for you.

## Most Brands Disappear for Category-Level and Feature-Specific Prompts

The citation data reveals a dominant connection between query type and consistency: there's a sharp drop-off in average mention rate when queries shift from branded (Comparison, Brand Evaluation) to unbranded (Category-Level, Feature-Specific). This is the Consistency Gap.

Query Type	Average Mention Rate
Comparison (branded)	98.7%
Brand Evaluation (branded)	98.5%
Category-Level (unbranded)	44.6%
Feature-Specific (unbranded)	42.0%

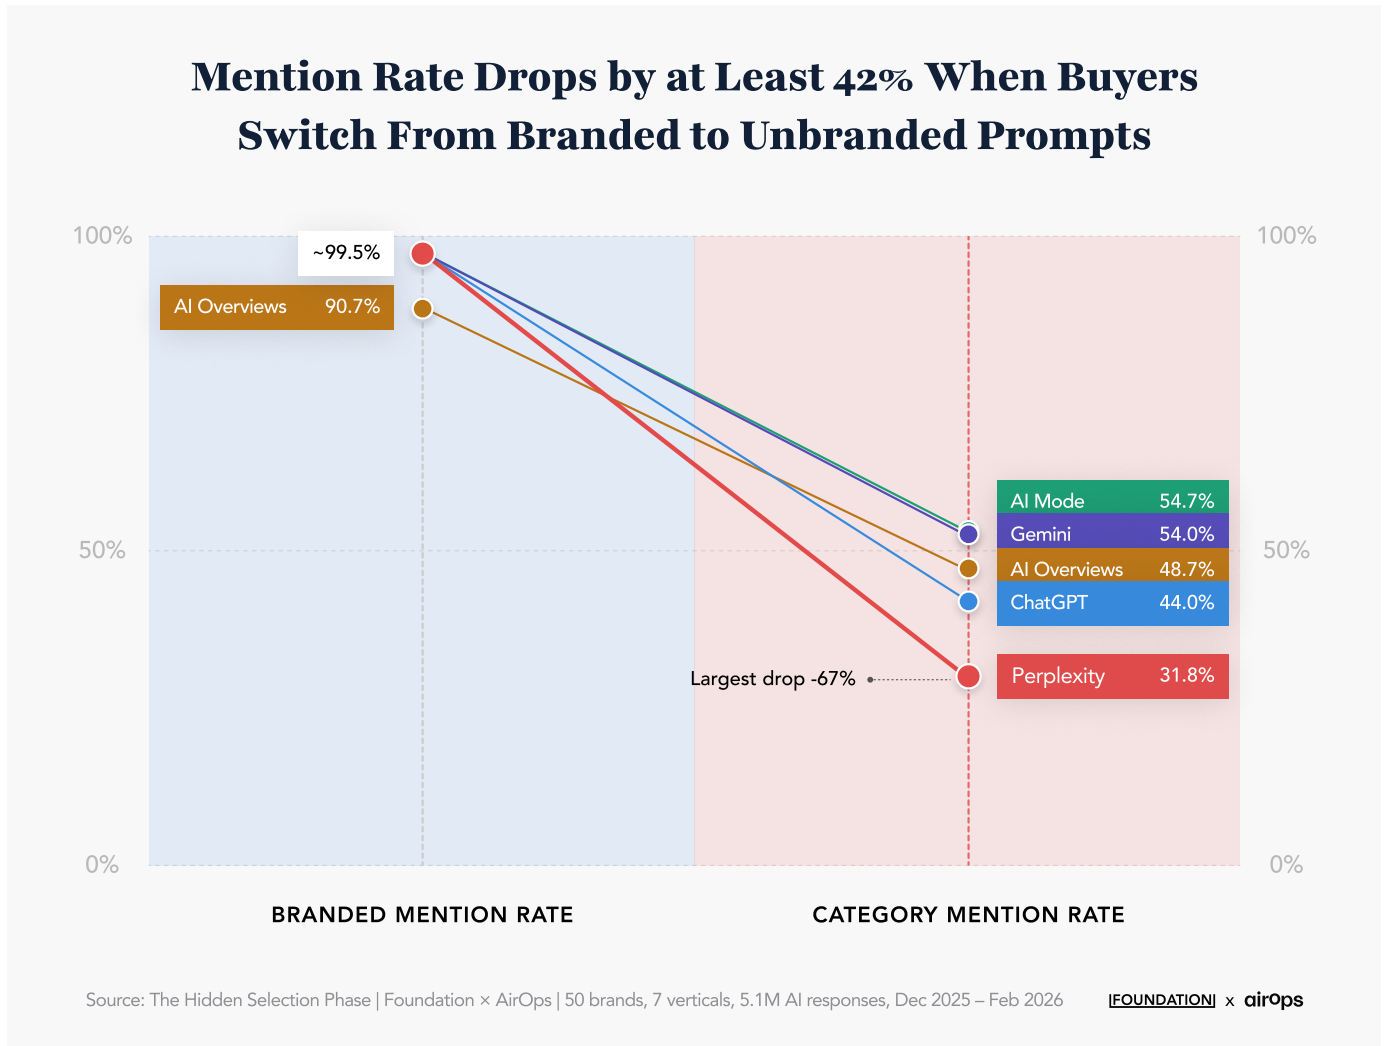
When a buyer asks about a brand by name, such as "How does [Brand] compare to [Competitor]?" or "Is [Brand] good for enterprise?", that brand appears in the response 98% or more of the time, which is expected. If you ask AI about a specific company, it will mention that company. It's also expected that any individual brand would appear less frequently in unbranded queries since the AI is drawing from a larger set of potential sources when no brand is named.

But the degree of variance is huge. When a buyer explores a category with prompts like "What is the best CRM for mid-market companies?" or "Which tools automate DevOps pipelines?", the average mention rate for any individual brand falls to 44.6% for category-level queries and 42.0% for feature-specific queries. That is a 55-plus percentage point drop.

The Consistency Gap is most impactful at the very moment AI shortlists are being formed. Category-level and feature-specific queries are the discovery moments, the prompts buyers use to form their initial impression of a market. Fewer than half of the brands appear in these responses. The shortlist is being assembled, and most brands are not in the room.

## Brand Visibility Varies By Query Type and AI Model

The Consistency Gap is not uniform across AI platforms. Some platforms are significantly more restrictive gateways for brand discovery than others.



Perplexity has the steepest Consistency Gap: 99% or higher on branded queries but just 31.8% on category-level and 26.3% on feature-specific. If a brand is not already well-established in Perplexity's sourcing, it is unlikely to appear in discovery queries.

Google's AI Mode and Gemini retain the highest unbranded mention rates at roughly 51–55%, making them the most accessible platforms for category discovery.

ChatGPT sits in the middle, and AI Overviews has a notably lower branded mention rate (90.7% vs. 99% + for all others), meaning it is more likely to skip mentioning a brand even when asked about it directly. The practical implication: a brand's AI visibility strategy cannot be platform-agnostic. Being discoverable on ChatGPT does not mean you are discoverable on Perplexity, and what works for AI Mode may not translate to AI Overviews.

## Brand Visibility Varies By Query Type and AI Model

The Consistency Gap is not uniform across AI platforms. Some platforms are significantly more restrictive gateways for brand discovery than others.

Vertical	Category- Level Mention Rate (Unbranded)	Feature-Specific Mention Rate (Unbranded)
Productivity	72.4%	65.1%
DevOps & Security	49.5%	48.7%
Marketing & Analytics	47.9%	47.7%
Fintech	43.8%	37.5%
Sales & Revenue	43.1%	37.0%
Healthcare & Wellness	40.3%	34.0%
Workforce & Comms	29.4%	25.9%

Productivity stands alone as the only vertical where the brands appear in the majority of category-level queries. Its 72.4% category mention rate is nearly double the next-best vertical. This likely reflects a market with well-known, established brands that AI models reference readily — names in the project management, note-taking, and collaboration space that have become near-synonymous with their categories.

At the other extreme, Workforce & Communications has the most fragmented discovery landscape. Brands in this vertical appear in fewer than one in three category-level queries and roughly one in four feature-specific queries.

Some of this vertical variation may reflect the breadth of each category. Verticals with more diverse product types, where brands solve different problems under the same umbrella — would naturally see lower per-brand mention rates on unbranded queries. Category-level mention rates, which ask about the broader solution space, are a cleaner measure of discovery visibility. Feature-specific rates should be interpreted with the caveat that not every brand in a vertical necessarily offers every feature queried.

## Only 11 of 50 Brands Show Up Consistently Across All Query Types

Of the 50 brands tracked, 11 qualified as “consistent”, maintaining a 70% or higher mention rate across three or more of the four query clusters. The remaining 39 were classified as “moderate.” No brand was classified as fully sporadic, although many had individual query types with mention rates well below 40%.

The distribution of consistent brands across verticals is uneven. Productivity accounts for four of the eleven, DevOps for three, and the remaining four are spread across Sales, Fintech, Marketing, and Workforce. Healthcare has none. Whether this reflects genuine differences in how AI models treat these industries or is an artifact of which brands qualified for the study and how broadly each vertical is defined is difficult to isolate with this dataset.

What the data does show clearly is that consistency is rare regardless of vertical — 39 of 50 brands fall short of the threshold — and that reaching it requires performing well not just on branded queries, where nearly every brand excels, but on the unbranded category and feature queries where most brands struggle.

## **Frequent Mentions Doesn't Mean You Control Your Story**

The data suggests that citation consistency does not guarantee strong citation ownership. The 11 consistent brands had a lower average citation ownership rate (26.95%) than the 39 moderate brands (35.55%).

A brand can be mentioned in AI responses constantly while the citation mix remains dominated by Reddit threads, review sites, and community content. Appearing frequently and being accurately represented are separate outcomes.

## **The Sources AI Uses Are Stable, But They Can Change**

Month-over-month citation retention across the 60-day window was high, ranging from 96.31% (Marketing & Analytics) to 100% (DevOps & Security, Fintech, Healthcare & Wellness, and Workforce & Communications). Mention retention followed a similar pattern, ranging from 96.43% to 99.69%.

This doesn't mean AI remembers what it said last month. Each response is generated fresh. What the retention data shows is source stability: the underlying content landscape that AI models draw from, the Reddit threads, help documentation, YouTube videos, and review site pages — doesn't change dramatically from month to month. The brands that appear consistently do so because they have already built a content presence in the places AI looks every time it generates a new answer.

This is an empowering finding, not a discouraging one. The source landscape is stable, but it is not static. Brands that invest in building presence across the sources AI draws from — comprehensive documentation, authentic community participation, video content, earned media — can change what AI says about them. Investing in the most citable content types on the channels LLMs defer to the most is a competitive advantage.


# Most B2B Brands Let Third Parties Write Their AI Story

The Attribution Gap tells you that 90% of your citations come from third parties. The Consistency Gap tells you that you may not appear at all in the queries that matter most. The Narrative Gap asks a different question entirely: when AI does cite those third-party sources about your brand, which sources are they and is the story they tell accurate, current, and aligned with how you would choose to position yourself?

## Reddit Controls 20% of All AI Citations in B2B SaaS

Reddit is the single largest external citation source in the dataset, accounting for 20.82% of top-50 external citations with 2.1 million citations across all 50 brands. It is the number-one external source in every vertical except Sales & Revenue, where LinkedIn narrowly edges it out (3.76% vs. 3.60%).

### Reddit Accounts for 20.8% of All External Citations



- reddit.com
- youtube.com
- Help/support docs
- Competitor SaaS
- Adjacent SaaS
- Uncategorized
- Other
- linkedin.com
- wikipedia.org
- Comparison/aff
- g2.com
- gartner.com

Share of top-50 external citation domains across 57.2M citations

<p>Branded queries</p> <p><b>28.0%</b></p> <p>Reddit share of external</p>	<p>Unbranded queries</p> <p><b>30.9%</b></p> <p>Reddit share of external</p>	<p>Verticals where #1</p> <p><b>6 of 7</b></p> <p>Reddit is top source</p>
----------------------------------------------------------------------------	------------------------------------------------------------------------------	----------------------------------------------------------------------------

Source: The Hidden Selection Phase | Foundation x AirOps | 50 brands, 7 verticals, 5.1M AI responses, Dec 2025 – Feb 2026 | FOUNDATION x airOps

Reddit's dominance becomes even more pronounced in discovery contexts. For unbranded queries, where buyers are forming first impressions through category and feature questions, 30.85% of external citations come from Reddit. For branded queries, the figure is 28.04%. When AI is building its answer to "What's the best [category] tool?", nearly one in three external sources is a Reddit thread.

This has significant implications. Reddit is user-generated, unmoderated by brands, and composed of individual opinions that may be outdated, unrepresentative, or actively hostile. Brands cannot control what Reddit says about them. They can, however, participate authentically in relevant communities, and the data suggests that this participation matters, not just because it gives brands control, but because it gives AI models more than one type of source to draw from.

## Query Type Changes Which Sources the AI Cites

Who shapes the narrative depends on whether the buyer asks about a brand by name or explores the category.

**In branded queries**, the external narrative is dominated by Reddit (28.0%), YouTube (19.6%), G2 (10.8%), and LinkedIn (8.5%). Together, these four sources account for 58% of the external sourcing. G2 is notably prominent here, when a buyer asks about a brand by name, review sites play a significant role in how AI frames the response.

**In unbranded queries**, the landscape shifts. Reddit rises to 30.9%. LinkedIn moves from fourth to second at 15.4%. YouTube drops slightly to 14.9%. Wikipedia comes in at 6.8%, and Medium at 5.9%. G2, which was 10.8% of the branded narrative, drops off the top 15 entirely for unbranded queries.

This shift has a critical strategic implication. G2 reviews shape how AI talks about you when someone already knows your name. They do not help you get discovered when a buyer is exploring the category. LinkedIn, conversely, plays a larger role in category-level narratives than in branded ones, and Wikipedia, which many B2B brands ignore entirely, is a meaningful source for the category-definition queries where first impressions form.

## Only 10% of Brands Actually Control Their Own AI Narrative

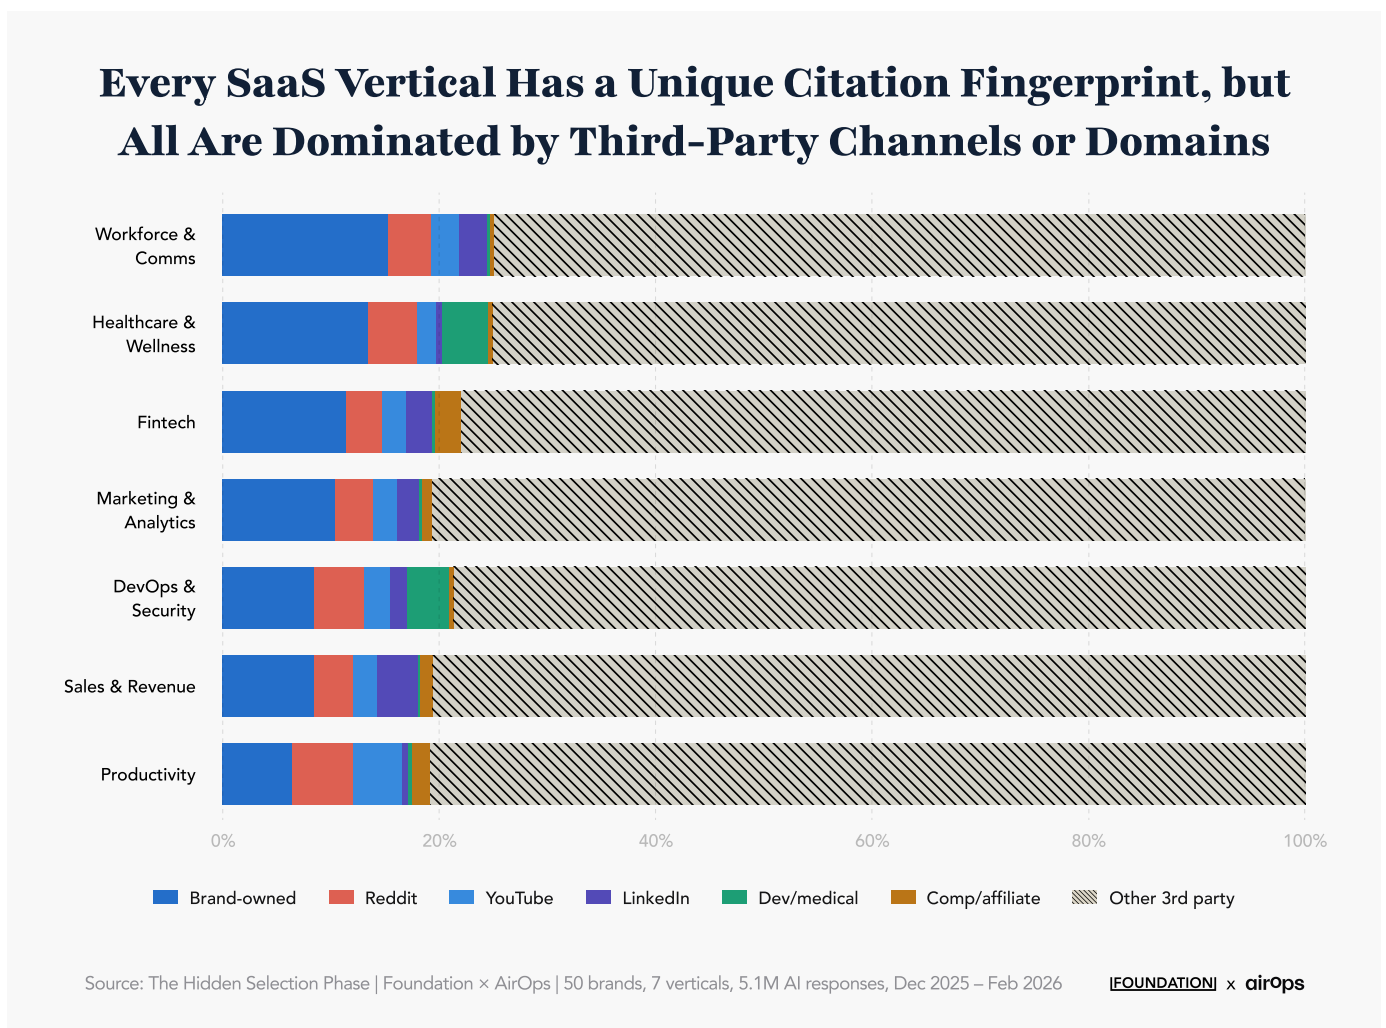
To assess narrative alignment, we examined the composition of external citations for each brand: specifically, which source types dominate, and whether brand-owned content is present in the mix alongside them. Brands were distributed across three tiers based on whether their owned content appeared consistently enough to compete with third-party sources in shaping the AI's framing:

- **High narrative alignment:** Brand-owned content present in 50%+ of responses, with owned sources appearing alongside community and third-party citations
- **Mixed alignment:** Brand-owned content present in 25–50% of responses; narrative shaped by a blend of owned and external sources
- **Third-party framed:** Brand-owned content present in fewer than 25% of responses; external sources set the framing with minimal owned counterweight

Tier	Threshold	Brand Count	% of total
High narrative alignment	50%+ response-level brand citation	5	10%
Mixed alignment	25–50%	27	54%
Third-party framed	Below 25%	18	36%

Only 10% of brands studied have high narrative alignment. More than a third have their story primarily shaped by third parties. The majority sit in the middle, and their owned content appears sometimes, but the narrative is a blend of brand messaging and third-party interpretation.

When we zoom in at the level of industry verticals, the extent of third-party influence begins to change.



**Fintech (11.43% brand owned):** NerdWallet (1.18%) and Bankrate (1.12%) are notable not just for their share of Fintech citations, but for what their presence reveals about the vertical. Comparison and affiliate sites do not appear in the top external citation sources for any of the other six verticals studied. The presence of multiple websites of this nature is a Fintech-specific phenomenon, and it suggests the category narrative in this vertical is being shaped by sites whose business model depends on actively ranking products against each other.

**DevOps & Security (8.46% brand owned):** Developer community platforms appear here in concentrations not seen elsewhere. GitHub (1.28%), Medium (1.85%), and dev.to all feature in the top external sources. Again, these are small percentages individually, but they represent a source category developer-generated technical content developer-generated technical content — that is distinctive to this vertical. Technical buyers evaluate tools through code examples, tutorials, and developer experience write-ups, and AI models reflect that evaluation pattern in their sourcing.

**Healthcare & Wellness (13.56% brand owned):** Medical sites with high Domain Authority create a distinctive layer of trust. PubMed (2.72%), Healthline (0.92%), and Mayo Clinic (0.64%) carry weight in AI responses that most brands cannot replicate at scale. These sites carry a type of institutional credibility that most B2B brands cannot produce themselves. The competitive dynamic here is different from other verticals: in addition to their direct competitors, brands are also up against major clinical and academic websites for authority in AI responses.

**Marketing & Analytics (10.40% brand owned):** Help documentation from brands in this vertical accounts for 1.45% of external citations. This is a higher share than in most other verticals, suggesting that brands investing in comprehensive support content are appearing in the citation mix more frequently. Integration platforms also play a distinct role: Zapier accounts for 0.83% of citations because it publishes workflow and integration guides that reference the tools it connects to. This is a pattern worth noting for brands in integration-heavy categories, even though the absolute numbers are modest.

**Sales & Revenue (8.43% brand owned):** This is the only vertical where LinkedIn (3.76%) outranks Reddit (3.60%) as the top external source. G2 (1.14%) and Salesforce.com (0.99%) also feature more prominently than in other verticals. The distinctive pattern here is the concentration of professional network and review platform citations. This suggests that thought leadership and executive visibility on LinkedIn, and active management of review profiles on G2, may matter more for sales tech brands than for brands in other verticals.

**Productivity (6.30% brand owned):** The lowest brand-owned citation rate of any vertical. At 4.60%, YouTube's citation share in Productivity is the highest recorded for any single source across the study, reflecting the visual, demo-heavy nature of how productivity tools are evaluated. Integration and ecosystem platforms (Zapier at 1.63%, Atlassian at 0.87%, Monday.com at 1.01%) also appear in higher concentration than elsewhere, which makes sense given that productivity tools are often evaluated based on what they connect to. The distinctive pattern: video content and ecosystem adjacency play a larger role in this vertical's citation mix than in any other.

**Workforce & Communications (15.18% brand owned):** The highest brand ownership rate overall, driven largely by robust help documentation. But this vertical also has the widest gap between branded and unbranded performance: 86.13% of branded responses include a brand-owned citation versus just 7.63% of unbranded responses — a 78.5-percentage-point gap. Brands here own their narrative when someone asks about them by name but are nearly absent as sources when buyers explore the category.

## Absence From AI Responses is a Much Greater Risk Than Negative Sentiment

AirOps applied automated sentiment classification to every brand mention in the dataset, which adds a dimension the original study framework did not anticipate.

AI responses are overwhelmingly either positive or neutral. Negative sentiment is rare, below 3.7% in every vertical-platform combination. The most negative pairing in the entire dataset, Sales & Revenue on Perplexity, reaches just 3.7%.

Based on this data, the risk of AI saying something negative about B2B brands is low. The most pressing risk is actually two-fold:

1. The AI does not mention the brand at all
2. It does mention the brand, but the narrative is misaligned or being shaped through third-party channels.


Sentiment also varies depending on the AI platform used. AI Overviews is the most positive platform in every single vertical, reaching 74.4% positive sentiment for Productivity brands. Gemini is the most critical, carrying the highest negative sentiment rate in four of seven verticals. AI Mode defaults to descriptive rather than evaluative framing, making it the most neutral overall. These differences suggest that brands should monitor not just whether they appear, but how each platform frames them.

# What Separates Brands AI Cites vs. Brands AI Ignores

The three gaps are not equally distributed. Some brands perform meaningfully better across attribution, consistency, and narrative control. Studying what separates them from the rest reveals actionable patterns.

## High Citation Ownership and Consistency is the Ideal State for AI Visibility

When we plot consistency score against citation ownership for all 50 brands, four distinct quadrants emerge. Each represents a different strategic position and a different set of problems to solve.



## High Consistency + High Ownership: The Ideal

These brands appear reliably across query types and maintain control of their narrative when they do.

- Healthcare-02 (73.25 consistency score)
- 77.84% citation ownership)
- Productivity-03 (80.11, 54.39%)
- MarTech-06 (57.88, 60.05%) exemplify this position.

They have built content ecosystems that AI models both find and cite, not only in branded queries but in the category and feature queries where discovery occurs.

## High Consistency + Low Ownership: Visible but Not in Control

These brands are mentioned everywhere, but AI cites third-party sources about them. Sales-01 has the starkest profile: an 84.88 consistency score paired with just 5.32% citation ownership. DevOps-02 (93.51, 21.65%) and Productivity-01 (89.91, 23.59%) show similar patterns. These are well-known brands with strong market presence, and AI mentions them constantly. But when it does, the citations point to Reddit threads, review sites, and competitor content rather than the brand's own domain. They have brand awareness in AI without narrative control.

## Low Consistency + High Ownership: Strong When Present, Often Absent

These brands own their story when AI mentions them but are missing from too many query types. Healthcare-05 (58.31, 51.26%), Fintech-03 (51.82, 48.65%), and Workforce-04 (57.45, 45.65%) fall into this category. Their content is authoritative enough for AI to cite it when relevant, but they lack the breadth needed to appear across the full range of buyer queries. The solution is usually content expansion rather than improving quality, covering more topics rather than better coverage of existing topics.

## Low Consistency + Low Ownership: The Double Gap

These brands are both inconsistently present and poorly cited from owned sources when they do appear. DevOps-08 (51.97, 40.55%) and Sales-07 (52.48, 38.98%) represent this quadrant. For these brands, the challenge is foundational. They need to build citation-worthy content and expand its coverage simultaneously.

## Support Documentation is a Strategic Moat for AI Visibility

The strongest correlation in the data linking owned content to citation performance comes from help documentation. Multiple brands' help and support subdomains appear in the top external citation sources for their vertical. In Workforce & Communications, one brand's help center accounts for 2.19% of all external citations in the vertical. In Sales & Revenue, two brands' help centers each account for more than 1% of vertical citations. In Marketing & Analytics, one brand's help center accounts for 1.45%.

These are effectively owned-narrative sources. The content lives on a subdomain and may be technically classified as external in citation tracking, but the brand controls every word. Help documentation creates citation surface area that AI models draw from across all query types, specially feature-specific queries where buyers ask about capabilities and integrations.

The implication is counterintuitive for many any B2B marketing teams: your help center may be a more important GEO asset than your blog. Comprehensive, well-structured documentation that answers the specific questions buyers ask AI is precisely the kind of content that earns persistent citations.

## Understand Reddit's Outsized Impact on AI Responses

**Reddit is the number-one external citation source in six of seven verticals.** Every brand in the dataset has Reddit citations in its AI responses, so Reddit presence alone is not a differentiator. What differentiates is whether a brand has owned content that competes alongside those Reddit citations, because the mix of sources AI draws from determines whose framing the buyer receives.

When AI cites your help documentation alongside a Reddit thread, the response blends your positioning with the community's perspective. The buyer gets your language, your feature descriptions, and your use cases, alongside whatever a Reddit user said. But when AI cites only Reddit, a G2 review, and a competitor's blog, the buyer gets everyone else's version of your story. In other words, you are present in the response, but the narrative belongs to someone else.

**Brands with 50% or higher citation ownership** all have substantial owned content in addition to their Reddit presence. They have not replaced Reddit but they have built enough authoritative content that AI models have brand-owned sources to cite alongside the community threads. The result is a blended narrative where the brand's voice is part of the mix rather than absent from it.

**For brands below the 50% threshold**, the narrative is disproportionately shaped by community content. The sentiment data shows that this doesn't necessarily mean the narrative is negative. AI responses skew positive regardless of source. But positive and accurate are not the same thing. What Reddit users said about your product three years ago may be what AI is telling buyers today, and that characterization may no longer reflect your current capabilities, positioning, or competitive advantages.

## Platform Differences Create Strategic Complexity

Not all AI platforms behave the same way, and the differences are large enough to matter for strategy.

Google's AI Mode generates an average of 20.05 citations per response, roughly three times more than ChatGPT's 6.12. Perplexity generates 10.20. As a result, the citation landscape on AI Mode is significantly more fragmented, with more third-party sources per response, increased competition for citation share, and a higher bar for owned-content visibility.

Perplexity, despite generating more citations per response than ChatGPT, has the lowest unbranded mention rate at 29.07%. It is the most restrictive gateway for brand discovery. If your brand is not already in Perplexity's sourcing for category queries, you are unlikely to appear in discovery moments.

AI Overviews, meanwhile, is the most brand-friendly platform for sentiment (68.09% positive across all verticals) but the most likely to skip mentioning a brand on branded queries (90.68% vs. 99%+ for all other platforms). It is simultaneously the kindest and the most selective.

**"Winning isn't about any single source. It's about stacking: your on-domain content plus your share of Reddit plus earned media plus review sites and so on. No single source wins alone; it's the aggregate. That's GEO."**

James Scherer, VP of Strategy, Foundation

# GEO Is Not SEO: Update Your Approach Accordingly

SEO optimizes for clicks to your site, while GEO optimizes for narrative presence inside AI responses. Rather than a rebranding of SEO, GEO represents a significant expansion of the objective and scope search visibility to include third party-platforms.

With SEO, the goal is to rank for a keyword so that a searcher clicks through to your domain. The traffic arrives, and your website does the work of converting it. With GEO, the goal is to be cited, mentioned, and accurately represented inside the AI's answer, because, by the time a buyer clicks through (if they click through at all), the evaluation is largely complete. The shortlist has been formed inside the LLM dialogue, and that shortlist is where 80% of purchase decisions are locked in.

This shift does not make SEO irrelevant. Organic search still drives the majority of digital traffic, and AI models draw heavily from content that ranks well. But it does mean that ranking is no longer sufficient. You can rank first for a keyword and still be absent from the AI's answer. You can have the best content on the web and still have your narrative told by Reddit.

## The 4 Things Every B2B Brand Should Do Right Now

### 1. Own Your Citations

The target: **40% or higher response-level brand citation rate**, based on the threshold analysis in this study. Brands above this level had average consistency scores of 73.8, compared to 64.1 for those below.

The starting point is an audit: where does AI currently get its information about your brand? Which third-party domains dominate? What content on your own domain is being cited versus ignored?

On-domain content alone accounts for roughly 10% of all citations in the dataset. It is necessary but insufficient. The path to 40% requires not just better owned content but also cultivated presence across the source types AI models actually cite: help documentation, technical guides, community participation, and earned media.

### 2. Be Consistently Present

The target: **70% or higher appearance rate across core query clusters**. Only 11 of 50 brands in this study achieved this. The critical gap is between branded queries (where virtually everyone appears) and unbranded category and feature queries (where the average drops to 43%). Consistency means being present when buyers are exploring, not just when they're validating. Platform fragmentation makes this harder. A brand that appears reliably on ChatGPT may be nearly invisible on Perplexity. Monitoring and strategy must be platform-specific.

### 3. Control Your Narrative

The target: owned content present in 50% or more of AI responses, competing alongside third-party sources rather than being absent from the citation mix entirely. Only 5 of the 50 brands studied achieved this. The rest have their narrative shaped predominantly by Reddit threads, review sites, and community content. Those sources may be positive in sentiment, but they're not necessarily aligned with your brand's current positioning, capabilities, or competitive framing. Narrative control is not just about sentiment but accuracy and alignment.

### 4. Treat AI as Infrastructure

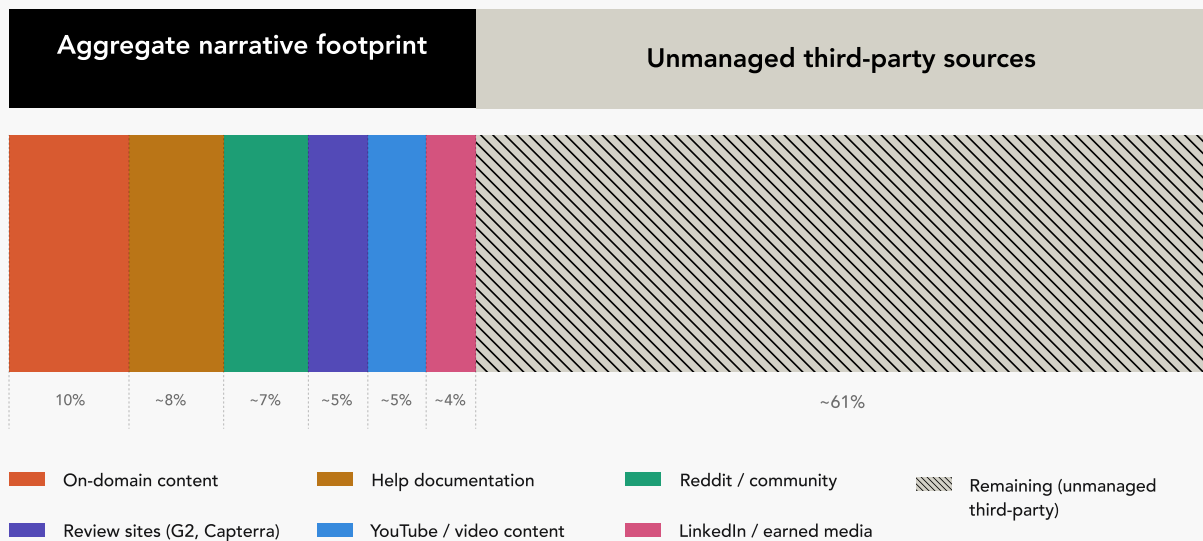
When 94% of buying groups use LLMs, 30% of buyers start with AI before Google, and LLM referral traffic converts at far higher rates than traditional search, AI is not a channel to experiment with. It is infrastructure. It is the environment in which buying decisions increasingly begin and often end.

The 96-100% citation retention rates in this study suggest that the competitive landscape in AI is hardening. The brands that invest now will build citation advantages that persist. The brands that wait will face an increasingly entrenched competitive set.

## No Single Source Wins AI Visibility. GEO is an Aggregate Game.

The most effective approach to GEO is not about dominating any single source. It is about accumulating presence across the source types AI models actually draw from.

### GEO is Additive and Third-Party Sources Make Up the Majority of Your Footprint



“ Your 5% on-domain plus your share of Reddit plus earned media plus review sites. No single source wins alone. That's GEO. ”

Your 10% on-domain citation share, combined with your presence in Reddit discussions, G2 reviews, YouTube content, LinkedIn thought leadership, help documentation, and earned media coverage, creates an aggregate narrative footprint. No single source will deliver 50% citation ownership, but stacking presence across all of them can.

This is what separates the brands in the top quadrant of the consistency-ownership matrix from the rest. They have not found a single hack or a single channel. They have built citation surface area across multiple source types, giving AI models a rich and diverse set of brand-aligned content to draw from regardless of the platform or query type.

## What Has Not Changed

It is worth noting what this research does not displace. 97% of buyers report prior experience with shortlisted vendors. Preliminary consensus within a buying committee still predicts 80% of wins. Brand reputation continues to outweigh feature comparisons in final decisions.

What has changed is where and how those perceptions form. They now form inside LLM responses during the Selection Phase. Every response is a brand touchpoint. For most B2B companies, it is an unmanaged one. The data in this report quantifies exactly how unmanaged it is, and provides a framework for changing that.

## The AI Selection Phase Is Already Happening. The Question Is Whether You're In It.

The data in this report is clear: 94% of B2B buying groups are using AI. Shortlists are being formed inside LLM conversations before a single sales call is booked. And right now, 90% of what AI says about most brands comes from sources those brands don't control. That gap is not permanent. It is measurable, and it is closable.

**Foundation is a B2B content marketing agency helping brands get found, cited, and accurately represented in AI responses.**

If you want to understand where your brand stands across the three gaps — attribution, consistency, and narrative — and build a strategy to close them, we'd like to talk.

[Let's Talk GEO Strategy](#)

*The Hidden Selection Phase is a Foundation × AirOps Research Report. © 2026 Foundation Inc. All rights reserved.*

# Appendix: Methodology Detail

This study was conducted in partnership with AirOps, whose AI visibility platform provided the underlying data infrastructure. The goal was to measure, at scale, how B2B brands appear in AI responses, and specifically to quantify the gap between brand-owned and third-party sourcing across verticals, platforms, and query types.

## Sample Size

We tracked 50 brands across seven B2B verticals over a 60-day period from December 22, 2025 to February 20, 2026. The seven verticals were selected based on the availability of qualifying brands with sufficient historical data and include:

- DevOps & Security
- Fintech
- Healthcare & Wellness
- Marketing & Analytics
- Productivity
- Sales & Revenue
- Workforce & Communications

Each vertical contains between six and eight brands, for a total of 50.

Responses were tracked across five AI platforms: ChatGPT, Gemini, Perplexity, Google AI Mode, and Google AI Overviews. The resulting dataset comprises more than 5.1 million AI responses and 57.2 million individual citations collected across these brands over a 60-day period.

Brand names are anonymized throughout this report. Each brand is referenced by its vertical and a numeric ID, such as “Fintech-03” or “Healthcare-07”. Any SaaS brand mentioned by name is the result of strong citation presence, not because of inclusion as a study participant.

## Prompt Structure

For each brand, 100 prompts were constructed across four query clusters and two intent types.

### 1. Unbranded Prompts

Sixty-five percent of prompts were unbranded, meaning the buyer does not name a specific vendor. These include 33 category-level prompts, such as “What is the best project management software for remote teams?”, and 32 feature-specific prompts, such as “Which tools integrate with Salesforce for pipeline reporting?”. These prompts represent discovery moments, when buyers find and evaluate solutions.

## 2. Branded Prompts

The remaining 35% were branded, where the buyer names a specific brand. These include 18 comparison prompts, like “How does [Brand] compare to [Competitor]?” and 17 brand evaluation prompts, such as “Is [Brand] good for enterprise?”.

These prompts represent validation moments, when buyers assess a brand they already know by name. Five thousand total prompts across 50 brands, run across five platforms, generated the 5.1 million response dataset.

## What We Measured

Every citation in every response was classified as either brand-owned (the brand’s tracked domain and subdomains) or external (everything else).

External citations were further categorized by source domain type, including:

- Review sites
- Community and UGC platforms
- Video and media
- Professional networks
- Competitor and adjacent tools
- Reference sites
- Comparison and affiliate sites
- News and media
- Brand help and support documentation.

## Citation Ownership

We divided our citation measurements into three categories based on whether they came from a branded domain or a third party source. This allows us to assess the scale of brand visibility (and absence) across what we call the The Three AI Visibility Gaps:

- **The Attribution Gap** measures the percentage of citations that come from the brand’s own domain versus third-party sources. This was calculated at both the citation level (each individual citation) and the response level (whether at least one citation in a given response pointed to a brand-owned domain).
- **The Consistency Gap** measures how reliably a brand appears across query types. A brand is classified as “consistent” if it maintains a 70% or higher mention rate across three or more of the four query clusters. “Moderate” indicates mixed performance (40-70% mention rate), while “sporadic” indicates sub-40% rates.
- Month-over-month citation retention was also tracked to measure the stability of AI citation sets over time for the same prompt.

- **The Narrative Gap** examines who else influences your brand's story inside AI responses. This is assessed through external domain distribution, which third-party source types dominate the citations, and through sentiment classification (positive, neutral, negative) applied by AirOps to every brand mention.

Throughout this report, we also measured citation ownership at two levels to illustrate the difference between showing up in a response and being a prominent source within that response:

**Response-level** asks a straightforward question: for a given AI response, did at least one citation point to the brand's own domain? This is a higher bar in one sense (it only counts yes/no per response) but a lower bar in another (the brand only needs one of potentially many citations to be its own).

**Citation-level** counts every individual citation in every AI response and asks: does this specific citation point to the brand's own domain or to a third party? Across 57.2 million citations, this gives us the broadest measure of source ownership. Both metrics matter. Response-level tells you how often your content shows up at all. Citation-level tells you how much of the overall sourcing you own.

**Brand selection:** 50 brands across seven verticals, selected based on a minimum of 60 days of consistent historical data within the December 22, 2025 – February 20, 2026 measurement window. Distribution of six to eight brands per vertical reflects qualifying brands within each industry during the study period.

**Prompt construction:** 100 prompts per brand. 33 category-level + 32 feature-specific (unbranded intent) + 18 comparison + 17 brand evaluation (branded intent). Prompts designed to represent realistic buyer queries across the discovery and validation phases.

**Data collection:** AirOps AI visibility platform, daily tracking cadence, across ChatGPT, Gemini, Perplexity, Google AI Mode, and Google AI Overviews.

**Citation classification:** Every citation classified as brand-owned (brand's tracked domain and subdomains) or external. External citations further categorized by domain type.

**Sentiment methodology:** AirOps automated sentiment tagging applied to each brand mention, classifying as positive, neutral, or negative.

**Anonymization:** Brand names are anonymized throughout. Each brand is referenced by vertical and numeric ID (e.g., Fintech-03). Separate internal key maintained.

## Limitations

LLM responses are non-deterministic. Daily variance in response content and citation sourcing is expected. The data represents patterns across the 60-day window, not guaranteed outputs for any single query.

Prompt volume is directional. No one has direct access to the actual distribution of real-world LLM queries. The 5,000 prompts were designed to represent realistic buyer behavior but may not perfectly reflect true query distribution.

Six to eight brands per vertical may not be fully representative of competitive landscapes that include dozens of players. Findings should be interpreted as directional patterns within the studied set.

Sentiment scoring is automated and may not capture nuanced or contextual sentiment shifts. The methodology should be evaluated in light of AirOps' specific classification approach.

This study measures AI visibility and citation patterns. It does not directly measure downstream buyer outcomes such as shortlist inclusion, win rates, or deal closures. Any correlations between three-gap performance and buyer outcomes are inferred from third-party buyer research rather than observed directly in this dataset.

Content-level metadata, such as word count, format, publish date, author credentials, was not available for cited URLs. Inferences about content characteristics are drawn from domain-level patterns and retention data, not direct content analysis.

## Sources

1. G2, 2025 Buyer Behavior Report. <https://learn.g2.com/2025-g2-buyer-behavior-report>
2. 6sense, 2025 Buyer Experience Report. <https://6sense.com/science-of-b2b/buyer-experience-report-2025/>
3. Ahrefs, ChatGPT Has 12% of Google's Search Volume but Google Sends 190x More Traffic to Websites. <https://ahrefs.com/blog/chatgpt-has-12-percent-of-googles-search-volume/>

# Acknowledgments

This report is the result of a genuine collaboration between Foundation and AirOps, and it would not exist without the following people.

**Dozie Anyaegbunam**, Creative Director at Foundation, developed the core thesis behind this report and the Three Gaps Model that structures everything in it. The framework is his.

**Oshen Davidson**, Head of AI Research and Product Intelligence at AirOps, made this report possible. Oshen collected, organized, and analyzed the citation, response, and prompt data across all 50 B2B brands — 5.1 million responses and 57.2 million citations that form the backbone of every finding here.

**Ethan Crump**, Marketing Lead at Foundation, wrote the report and directed the graphic design process from brief to final execution.

**Nicole Menendez**, Senior Editor at Foundation, provided strategic and editorial review throughout, sharpening both the arguments and the prose.

**Pedro M**, Editor at Foundation, provided additional editorial review and ensured the final copy was clean and consistent across every section.

**Rizki Agus**, Graphic Designer at Foundation, created the data visualizations and designed the report you're reading.

---

*Thank you to everyone at Foundation and AirOps who contributed feedback, reviewed drafts, and helped bring this research to life.*

# The AI selection phase is already happening.

## *Are you in it?*

How Foundation helps B2B brands close the three gaps and win inside AI-driven search.

94% of B2B buying groups now use LLMs during research, shortlists lock in before a single sales call, and 90% of what AI says about most brands comes from sources those brands don't control. That gap is measurable, closable, and exactly what Generative Engine Optimization is built to solve.

### **01**

#### **Audit**

Map your citation footprint across the three gaps.

### **02**

#### **Build**

Create the content AI actually cites in your vertical.

### **03**

#### **Distribute**

Show up across every surface that shapes your narrative.

## **Ready to see where you stand?**

Book a free GEO strategy session. We'll map your brand across attribution, consistency, and narrative, and show you where the fastest wins are hiding.

[Let's Talk GEO Strategy](#)

[foundationinc.co/services/generative-engine-optimization-agency](https://foundationinc.co/services/generative-engine-optimization-agency)

The Hidden Selection Phase is a Foundation × AirOps Research Report.

© 2026 Foundation Inc. All rights reserved.