WORLDCOM
Public Relations Group
The Local Advantage. Worldwide.

WHERE TRUST LIVES



SOURCES FOR CREDIBILITY IN LOCAL MARKETS

JUNE 2026

RESEARCH PROJECT WITH
SEMRUSH
An Adobe Company



What The Shift to the Authority Economy Means

With the rise of ChatGPT and Artificial Intelligence (AI), it is easy to understand that there is a shift in the search market. Companies have been feeling the impact on their performance with declining numbers on website analytics.

In the past, when someone queried Google about a topic, brand or market, the response provided websites pages that were ranked based on 'matching' the query users were searching for. Now, searches in AI tools provide the response as an answer, not just the 'source' to potentially find the answer.

However, the shift to AI responses hides another level of complexity. AI tools use Large Language Models (LLMs) to create responses for users. An LLM will frequently use a source as a grounding reference to grab a bit of information and begin constructing an answer. However, that source might never be surfaced as a source to the user.

Through the data in the report, we will see when brands are cited. That does not mean they will gain visibility or mentions in AI responses. However, it does contribute to a brand's overall authority. In the new Authority Economy, it is important to realize that 75-85% of brand mentions come from external sources [1].

The report implies that if a domain is cited, brands gain visibility, but that's not always true. It's worth noting that citation frequency in the data reflects how often AI systems draw from those sources. It does not correlate to how often users actually see them mentioned as sources. For PR teams, that's a critical distinction.

Your brand's reputation is now shaped not just by what you publish as part of a media push, but by what the surrounding source ecosystem says about you. And that ecosystem looks very different than the traditional media or search landscape.

The conversation a brand is having in one place needs not only to be consistent across the channels it uses, but among the external sources and social communities where those brands are mentioned.

Experts talk about how 'authentic' storytelling will help win in communications. However, brand-generated content and user-generated content about that brand are both used to create perception and authority in AI engines.



What's In This Report

From the study data, we examined the development of the Authority Economy, how that plays out through global and local citation data and ways to help build an AI visibility strategy.

- **The Authority Economy** — An introduction to how AI answers have changed the relationship between brand-published content and external source ecosystems, and why 75–85% of brand mentions in AI answers now come from sources outside the brand's own channels.

- **Participatory Platforms as the Answer Supply Chain** — An analysis of the worldwide top 15 cited domains showing that six of the fifteen are social, video, community, or professional platforms. The data reframes these channels not as distribution tools but as credibility inputs into AI-generated answers.

- **Engine-by-Engine Profiles** — A breakdown of how AI tools each behave differently as citation environments, including how top source preferences shifted across the 34-week study period for LinkedIn, Wikipedia, and Reddit.

- **How Local Context Changes the Source Set** —Local market citation examples showing how government domains, local publishers, commerce sites, and community environments enter and lead the citation mix when prompts are market-specific. This section directly addresses why global platforms alone are not sufficient for in-market AI visibility.

- **A Six-Layer Evidence Framework** — A practical model mapping the source types of AI systems draw with the communications role each play. These sources include community and discussion, video and demonstration, professional credibility, reference and research, editorial and business media, and local authority.

- **Six Workstreams for Building AI Citation Credibility** — Actionable guidance on source mapping, local evidence architecture, community intelligence, expert and professional visibility, reference accuracy, and ongoing citation monitoring.

The report closes with a clear argument: there is no separate AI credibility strategy. The disciplines that build genuine reputation — local expertise, subject matter authority, authentic community participation, credible third-party coverage — are exactly what AI citation environments reward.



Participatory Platforms Are Fuel for the Answer Supply Chain

To find the sources that were 'trusted' and cited the most often, we partnered with *Semrush* to do an *LLM Citations Study*[2]. This research provided prompt-level citation data across 34 weekly reporting periods from October 2025 through May 2026. The research covered ChatGPT Search, Gemini, Google AI Mode, and Perplexity not only from an aggregated worldwide perspective, but also did a deep dive across eight global markets including Argentina, Brazil, Germany, Indonesia, Singapore, South Africa, the United Arab Emirates, and the United States.

From the study data, there is clear evidence that AI often cites online communities, content creators, authoritative publishers, and institutions. It shows a heavy lean into these external sources not only on a global scale, but on an individual country level as well.

The most striking finding from the worldwide May 2026 data is that six of the top 15 cited domains are social, video, community, or professional platforms. With Reddit, YouTube, X, LinkedIn, Facebook, and Instagram in those positions, these platforms are not fringe citations. They serve as primary signals that build authority.

Prompt Share for Citations

31.5%	29.1%	19.0%	18.2%
Reddit — worldwide prompt share, May 2026	YouTube — worldwide prompt share, May 2026	X — worldwide prompt share, May 2026	LinkedIn — worldwide prompt share, May 2026

Prompt share = % of analyzed prompts in which each domain was cited. Shares are not mutually exclusive; one answer can cite multiple domains. Source: Semrush LLM Citations Study, May 2026 Worldwide combined view.

This reframes a common misconception. Social channels are often treated as distribution tools that are used to push messages outward. The data shows the conversational nature of social channels also makes them data input environments.

AI systems are reading what communities say (Reddit, X), what creators demonstrate (YouTube), and what professionals discuss (LinkedIn). Those communities are generating content that is cited as evidence in generated answers.

For public relations and communications organizations already using social channels for local impact and community engagement, this is a direct validation of that approach. Authentic participation in communities, expert commentary on LinkedIn, and video-based explanation on YouTube are not just reach tactics. They are credibility inputs into the AI answer layer.



Worldwide Top 15 Citation Domains — May 2026

#	Domain	Prompts Citing	Prompt Share	Source Type
1	**reddit.com**	**175,269**	**31.5%**	Community / Discussion
2	**youtube.com**	**161,978**	**29.1%**	Video / Platform
3	**x.com**	**105,501**	**19.0%**	Social / Discussion
4	**linkedin.com**	**101,178**	**18.2%**	Professional Network
5	**facebook.com**	**43,080**	**7.7%**	Social / Community
6	wikipedia.org	41,034	7.4%	Reference
7	google.com	36,909	6.6%	Platform
8	nih.gov	36,233	6.5%	Institutional / Health
9	medium.com	33,441	6.0%	Editorial
10	forbes.com	27,141	4.9%	Business Media
11	**instagram.com**	**26,738**	**4.8%**	Social / Visual
12	arxiv.org	25,431	4.6%	Research
13	microsoft.com	23,301	4.2%	Brand / Platform
14	researchgate.net	21,889	3.9%	Research
15	sciencedirect.com	21,187	3.8%	Academic / Research

Source: Semrush LLM Citations Study — Charts.xlsx, Top LLM Citations Combined tab; May 2026 Worldwide aggregate The rows highlighted in yellow shows the participatory platforms in the top 15. Prompt share reflects % of analyzed prompts citing each domain. However, the shares are not mutually exclusive.

The strength of the prompt share among these social communities shows that AI is not simply pointing people to official owned sources. It is pulling evidence for responses from community driven sources.

When the perception and authority for a brand is strongly influenced by the consumers' conversation about it, that makes authentic communications more important, not less.



Not All Engines Created Equal

The worldwide blended ranking is a useful headline, but it conceals a critical operational reality: each AI engine has its own audience and citation behavior. In the past, mastering Google algorithms meant you could have steady rankings in other engines. The AI engine world is far more complex.

A strategy built around the preference of one AI engine will potentially underserve the others. Brands need engine-specific visibility planning. This is not a technical observation. It is a strategic communications requirement.

The implication is that AI visibility requires a distributed and localized credibility strategy, not a single SEO playbook applied uniformly across platforms. However, creating a strategy that will surface in every AI tool is not a task you can accomplish without endless time and resources.

So where should you concentrate your effort to get the most return? Currently, the largest two sources for AI search referrals are ChatGPT and Google. Google AI Overview now has 2.5 billion active monthly users.[3]

In February, ChatGPT had 900 million weekly active users, 50 million consumer subscribers and 9 million business subscriptions.[4] Although the other engines are important, these two AI tools will have the biggest impact on performance and return on investment.



Is there a Universal Hierarchy for AI?

Even if I filter my targets to only ChatGPT and Google AI, these tools have different, and ever evolving, behaviors. It is important to understand the mix of sources that are used for citation.

Yes, each engine has a specific persona that can address different modes of communication and will need to be fed content and information in different ways for it to be cited and mentioned.

Engine-by-Engine Profile — May 2026 Worldwide

Engine	May 2026 Profile	Communications Action
ChatGPT Search	Community + reference mix	Link authority with community discussion
Gemini	Lower citation frequency; X, Reddit, Medium	Interpret with coverage context — compare carefully
Google AI Mode	Video and platform led: YouTube, LinkedIn, Reddit, social	Prioritize video, social proof, professional profiles
Perplexity	Professional and source-rich: YouTube, LinkedIn, X, NIH	Build research-grade, professional evidence layers

Source: Semrush LLM Citations Study, May 2026 Worldwide split-by-LLM tables and Updated Analysis. Note: Gemini cites in approximately 61.8% of prompts vs. ~95–96% for other providers; direct comparisons require citation-coverage adjustment.

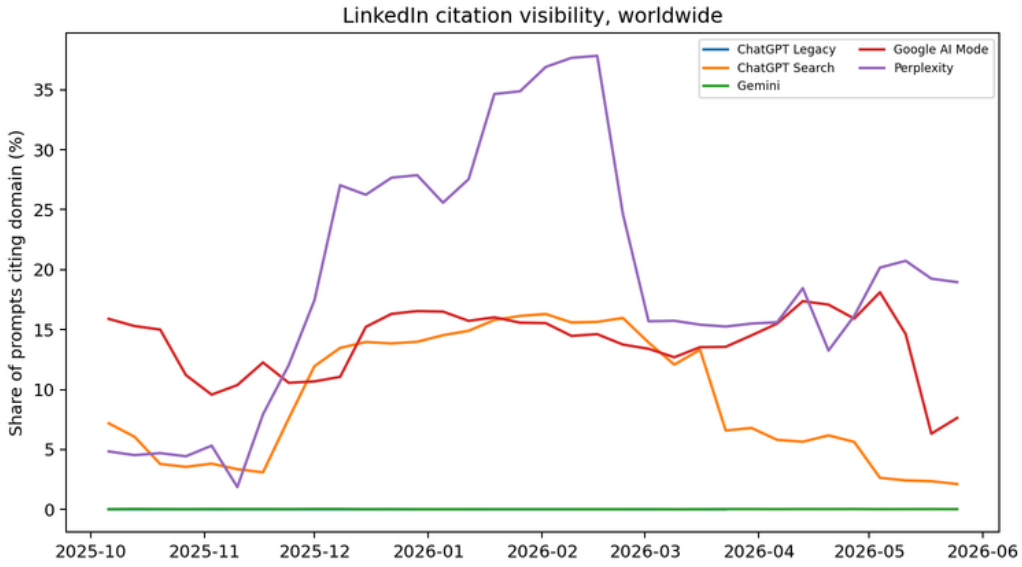
Even with the diversity of sources, there were some common sources found *Semrush* LLM *Citations Study* including LinkedIn, Wikipedia, and Reddit. To get a better insight into behavior and visibility, we looked at the performance and nuance of those three platforms across engines.

One of the things that is seen in the data for each of these sources below is that there can be dramatic shifts in citations. For example, Perplexity was heavily citing LinkedIn for the first half of the reporting period, but it charted a dramatic decline in February 2026 that didn't recover by the ending period in June. Wikipedia suffered the same fate in ChatGPT with a decline at the end of the data period that did not recover.

With these dramatic shifts, it is important that any plan remains fluid with sources and content as the AI platforms and methodologies mature and evolve.

LinkedIn

The AI visibility results for LinkedIn are highly platform dependent. It weakens materially in Google AI Mode and ChatGPT Search. However, there was a major shift in Perplexity as a trusted source by the end of the data period. Brands should therefore avoid treating professional-network visibility as a universal AI tactic.

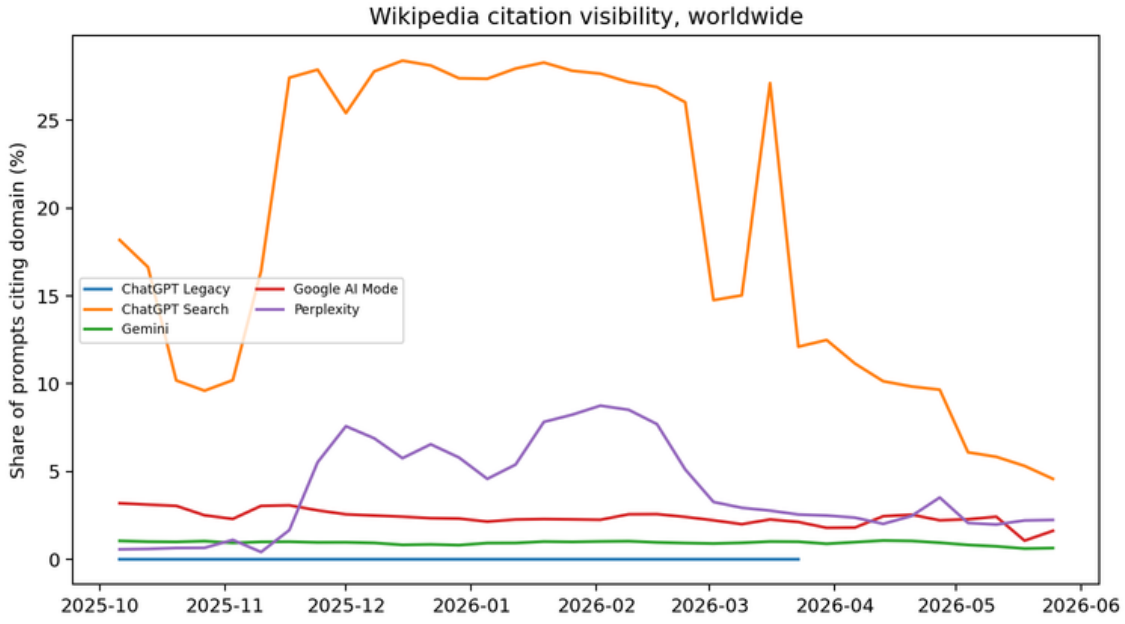


Provider	Period avg.	First week	Last week	Change
ChatGPT Search	9.5%	7.2%	2.1%	-5.1 pp
Gemini	0.0%	0.0%	0.0%	+0.0 pp
Google AI Mode	14.0%	15.9%	7.6%	-8.3 pp
Perplexity	19.8%	4.8%	19.0%	+14.1 pp

Wikipedia

Wikipedia remains the defining source for ChatGPT Search for most of the period, yet its final week share falls sharply. That shows that even Wikipedia shows inconsistency in AI results. It also declines in Google AI Mode and slightly in Gemini, while increasing in Perplexity. The pattern suggests that Wikipedia's role is not disappearing, but it is being redistributed across engines and over time.

This rise and fall pattern is something that Wikipedia experienced for years throughout traditional search when it was highlighted and removed as a trustworthy source in Google.

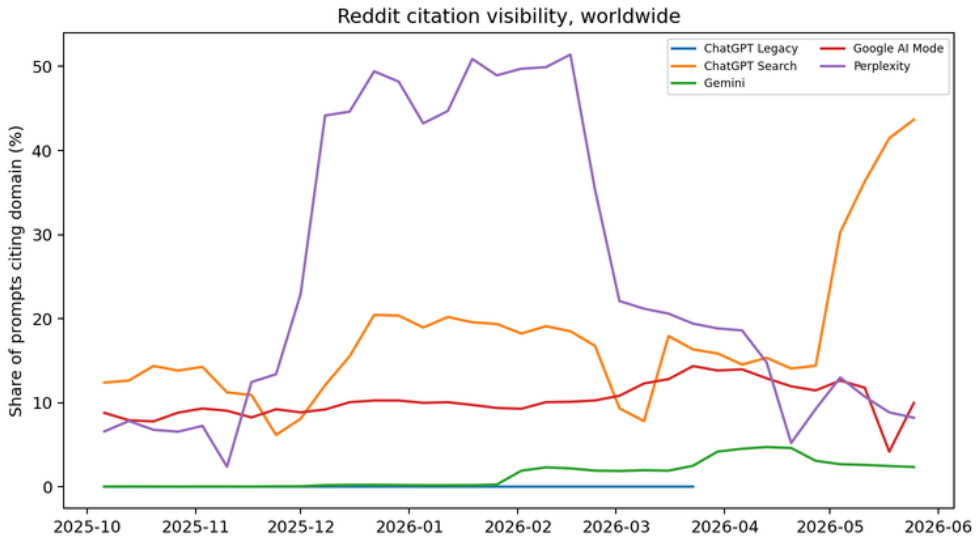


Provider	Period avg.	First week	Last week	Change
ChatGPT Search	19.2%	18.2%	4.6%	-13.6 pp
Gemini	0.9%	1.1%	0.6%	-0.4 pp
Google AI Mode	2.4%	3.2%	1.6%	-1.6 pp
Perplexity	4.1%	0.6%	2.2%	+1.7 pp

Reddit

Reddit is the strongest cross-engine winner, but for different reasons. However, there is a more erratic nature to its performance. It is structurally dominant in Perplexity, but it only experienced a peak for four months when it rapidly declines and dropped to previously nominal levels. However, it shows more consistency in Google AI Mode and shows a dramatic rise in ChatGPT Search by the final week.

The ChatGPT Search jump is large enough to warrant validation against product or ranking changes. It given some of the other inconsistent performances, it begs the question of whether the trend will be sustained. However, the broader conclusion is robust: community discussion on Reddit has become a core source type in AI answers.



Provider	Period avg.	First week	Last week	Change
ChatGPT Search	17.4%	12.4%	43.7%	+31.3 pp
Gemini	1.4%	0.0%	2.4%	+2.3 pp
Google AI Mode	10.3%	8.8%	10.0%	+1.2 pp
Perplexity	25.8%	6.6%	8.2%	+1.6 pp



Breaking Down the AI Tool Targets

There is a lot to digest with the variations of visibility across tools and the changes over time. A practical read for communications teams:

- **ChatGPT Search** shows a community-plus-reference profile. Strong performance requires linking community credibility (Reddit, X) with authoritative backing for sources with stronger vetting mechanisms (Wikipedia, NIH, arXiv).
- **Google AI Mode** is heavily video and platform led. YouTube, LinkedIn, Reddit, X, Facebook, and Instagram are all present in the May top 10. Video assets, social proof, and professional profiles structured for AI readability are the priority.
- **Perplexity** shows a strong professional and source-rich profile, with YouTube and LinkedIn ahead of Reddit in May. Research-grade, professional evidence layers carry significant weight.
- **Gemini** cites less often than overall. It has around 62% of prompts vs. 95%+ for other providers. Interpret Gemini data with care and always pair prompt share with citation coverage context.

Based on the sources for these engines, it indicates Reddit, LinkedIn and Wikipedia are strong drivers. If there was a magic combination, the data would point to ChatGPT and Reddit. One thing to consider is the significant shifts in performance. If there is no consistency with one AI source, what does that mean? It means that planning by using only large global platforms can create inconsistency in your results.

With the inevitable shifts in performance and visibility, brands need to use a nuanced approach to help build trust with consumers and authority with AI engines. People do not search on a global scale. They use their phones and sit behind computers in a specific location on the globe. It is important to remember that authority has a global impact, but visibility and citations happen on a local level in AI engines.



How Local Context Changes the AI Engine Sources

You might see the aggregate impact of Reddit, YouTube, X, and LinkedIn and lean hard into those sources for global coverage. However, the strategic differentiation is knowing local sources are not secondary. For market-specific prompts, they can be primary citation signals.

Part of high visibility is knowing which local sources carry credibility in a specific market. The stories that emerge in AI searches will be determined by how those local sources interact with global platforms.

Having a deep understanding of authoritative sources in a local market is exactly where a global network of local communications specialists has a defensible advantage. Local expertise is language-specific, category-specific, and culturally specific. It cannot be centralized or templated. It must be built in-market.

Local Citation Examples by Market — Final Week Available Data

Market	Local Citation Examples	What It Signals
Brazil	gov.br, globo.com, serasa.com.br	Government, major publishers, consumer finance appear alongside global social platforms
Singapore	ema.gov.sg, sureclean.com.sg, coreconceptsphysio.sg	Local service providers and government authorities are highly visible for market-specific prompts
UAE	shory.com, policybazaar.ae, falcontrackers.com	Commercial comparison, insurance, and technology sources appear with LinkedIn and Reddit
Germany	autoscout24.de, frankenschaum.de	Local commerce and specialist sources coexist with global platforms
United States	nih.gov, forbes.com, reddit.com	Deep and diversified: community, social, health authority, and business media all visible
Argentina	silvanadato.com, samsung.com, tiktok.com	Platform visibility mixes with local and category-specific sites

Source: Semrush LLM Citations Study — Charts.xlsx, weekly_raw tab. Domain examples are final-week top-domain signals across available providers. These are illustrative citation examples, not consumer trust rankings.



Final-week data from available markets shows that government domains, local publishers, commerce sites, insurance comparison platforms, service businesses, and market-specific communities enter, and sometimes lead, the citation mix alongside global platforms.

When looking at sources by country, the global impact of volume from U.S. sources begins to be filtered. You need to consider that 18.44% of ChatGPT's traffic comes from the United States according to Similarweb.[5] That means U.S. citations take on an outsized roll in a global overview.

Based on the data by country, there are specific websites that were proving to be sources during the last week of the data draws. These are broken out for reference. However, these sources should be considered with the same views as the inconsistencies in Reddit and ChatGPT over time.

Argentina

YouTube, Instagram, Reddit, TikTok, and local commerce or publisher domains are more visible than in the global average, reflecting a highly platform-driven citation environment.

Most persistent final-week domains across providers: samsung.com, reddit.com, youtube.com, silvanadato.com, esteticas.com.ar, instagram.com.

Brazil

Social, video, and community platforms are especially prominent, alongside local publishers and government domains. Google AI Mode shows a particularly strong social-platform signature.

Most persistent final-week domains across providers: youtube.com, serasa.com.br, google.com, gov.br, reddit.com, globo.com.

Germany

Global platforms remain dominant, but reference and institutional domains carry more weight. German-language publishers and local commerce or public-information sources enter the competitive set depending on provider.

Most persistent final-week domains across providers: microsoft.com, reddit.com, x.com, autoscout24.de, frankenschaum.de, youtube.com.



Indonesia

YouTube and social platforms are important, but local news, commerce, and government sources also shape answers. Lower prompt volumes in some provider-week combinations increase volatility.

Most persistent final-week domains across providers: nike.com, reddit.com, youtube.com, alibaba.com, runrepeat.com, runnersworld.com.

Singapore

A highly international, technology- and business-oriented source mix. Global professionals, research, and platform domains dominate more than local consumer portals.

Most persistent final-week domains across providers: sureclean.com.sg, ema.gov.sg, helpling.com.sg, coreconceptsphysio.sg, reddit.com, youtube.com.

South Africa

The citation landscape heavily leans to global and local providers of rental cars. However, Google.com and Reddit.com are present in the top citation group and remain important sources.

Most persistent final-week domains across providers: pacecarrental.co.za, avis.co.za, hertz.co.za, sani.co.za, budget.co.za, kenings.co.za.

United Arab Emirates

An internationally oriented citation market with strong technology, professional, travel, commerce, and government-source representation. Local authority sites can become prominent in market-specific prompts.

Most persistent final-week domains across providers: falcontrackers.com, shory.com, policybazaar.ae, lookinsure.com, youtube.com, linkedin.com.

United States

The deepest and most diversified citation market. Reddit, YouTube, X, LinkedIn, NIH, Forbes, Medium, and specialist review or finance sites all matter. Scientific and health authority is especially visible.

Most persistent final-week domains across providers: reddit.com, x.com, nih.gov, youtube.com, google.com, forbes.com.



From Message Distribution to Evidence Distribution

The strategic answer is not to chase every cited domain. It is to build a distributed credibility system matched to the type of content and authority each of these provides to AI Engines.

To understand how a source will contribute to a brand's perception, there are six evidence layers that AI systems draw from.

Evidence Layer	Typical Sources	Communications Role
Community & Discussion	Reddit, X, Facebook, Instagram	Where people ask questions, compare products, and validate claims in plain language
Video & Demonstration	YouTube	Tutorials, expert walkthroughs, reviews, and proof that can be cited and summarized
Professional Credibility	LinkedIn, G2, Gartner	Executive, practitioner, and B2B proof points — especially powerful in Perplexity and Google AI Mode
Reference & Research	Wikipedia, NIH, arXiv, ScienceDirect	Supports factual, health, technical, and regulated claims with high-authority backing
Editorial & Business Media	Forbes, Medium, TechRadar	Third-party validation, expert framing, and market context from trusted publishers
Local Authority & Services	gov.br, ema.gov.sg, local publishers	In-market credibility that global content cannot replicate — where local expertise wins

Source: Source-type interpretation based on Semrush cited-domain data; categories are analyst-defined for communications planning purposes.

The shift in frame is this: conventional campaign planning asks "what message are we publishing and where?" AI-era communications planning asks a harder question: "what evidence will AI systems and audiences find when they look for confirmation?"

A well-rounded program will not just target community and discussion, but it will also engage with channels that provide professional credibility, reference and community discussion.

It also means having a presence across the citation domains that are relevant to your brand. However, that content and conversation must be authentic because AI citations often surface places where audiences are already asking practical questions. A vague brand claim is less useful than a credible answer, a local proof point, a visible expert, a reliable third-party article, or a community discussion that validates the claim.



Five AI Citation Takeaways

1. Trust is becoming machine-visible through citations. But citations are not the same as trust. There is a new reputation surfacing that communications teams must understand and plan for.
2. AI answer systems reward source ecosystems, not isolated brand claims. Credibility must be distributed across the environments AI systems use for evidence.
3. Community, video, professional networks, publishers, institutions, and local authorities now sit in the same evidence chain. Targeting a single channel is sufficient.
4. Every engine behaves differently. A blended average hides the real activation challenge. Strategy must be segmented by engine, not applied uniformly.
5. Global brands must have localized approaches to ensure they are represented consistently and accurately in specific markets. Local expertise is the difference between global messaging and credible in-market evidence. This is where a global network of local specialists has an irreplaceable advantage.



Six Practical Workstreams for Brands

The data in this report points to a shift in authority that is already underway. AI answer systems are not reading brand websites and deciding whether they are credible. They are scanning a wider source ecosystem and assembling answers from what they find there.

The question communications teams need to ask themselves now is *what evidence will an AI system and our target audience find when they look for confirmation of our claims?*

Due to the nature of AI, the path to that answer is complex because the answers vary by engine, shift by market, but ultimately reward distributed credibility over centralized messaging.

The other challenge is that in many markets, for many topics, the online ecosystem around a brand is thin, inconsistent, or dominated by sources the brand has no relationship with.

The good news is that the same disciplines that build genuine reputation are exactly what AI citation environments reward. Demonstrating local expertise, subject matter expertise, credible third-party coverage, authentic community participation and accurate reference assets will drive positive visibility.

There is no separate AI credibility strategy. There is just a better version compared to the one a good communication strategy has always required. What follows is a practical framework for building and maintaining it.

Know your citation landscape before you try to change it.

Source mapping: Map cited domains by market, topic, and provider. Separate global platforms from local authorities.

The starting point is understanding which domains AI systems are citing when users ask questions relevant to your brand, category, or market. This is not the same as knowing which publications cover you. A brand may have strong editorial coverage but low AI citation visibility because the engines are drawing from community platforms, reference sites, or local authorities that the brand has never prioritized.

Effective source mapping means running prompt-level checks by market, by topic area, and by engine. The results will differ significantly across all three dimensions. What Google AI Mode cites for a health-related query in Germany will not look like what Perplexity cites for the same query in the United States. Using tools like *Semrush AI Visibility Toolkit*, helps to separating results to show local authorities. That will reveal where your brand has a strong citation presence and has genuine coverage gaps.



Use source mapping as a planning input, not a one-time audit. Citation environments shift as platforms evolve, as communities grow, and as AI providers update their models. A quarterly cadence is a reasonable starting discipline.

Global claims need local proof points.

Local evidence architecture: Create in-market proof points such as local case studies, local-language pages, expert quotes, partner content, and credible third-party references.

One of the clearest findings in this research is that citation source sets differ materially by geographic market. Global platforms appear everywhere, but local publishers, government agencies, service providers, and community environments enter the mix in ways that vary by geography, language, and category. A brand that relies entirely on centrally produced English-language content is leaving significant local citation territory unaddressed.

Building local evidence architecture means creating in-market proof points that give AI systems and audiences something verifiable to find. Using content like local case studies that demonstrate impact and cultural understanding in a specific market, local-language pages and resources that serve genuine user needs, expert quotes and perspectives from credible in-market voices, partner content that broadens the ecosystem of third-party references, and relationships with credible local institutions and publications.

This is not a translation exercise. Translating a global press release into another language does not create local evidence. It creates local text. The distinction matters: AI systems cite don’t just give answers because they are in another language. They are giving a response to something that fits with the local context and culture. Local evidence architecture is about building the assets that thoughtfully answer questions local audiences and AI systems are asking in-market.

Listen before you publish.

This should not be a new rule in the social sphere. However, it applies to any AI content strategy. With social, community, or professional platforms having a dominant position in AI, conversations from users are not a peripheral signal. They are central to it.

The practical implication is that community intelligence should precede content strategy, not follow it. Monitoring Reddit threads, YouTube comments, X conversations, LinkedIn discussions, and local community environments reveals the actual language audiences use to describe their problems, the objections and concerns that recur, the evidence gaps where



good information is absent, and the questions that AI systems are likely to surface when users ask about your space.

This intelligence shapes better content, better community participation, and better understanding of where your brand's credibility is strongest and weakest in the eyes of the communities that AI systems cite most frequently. It also serves as an early warning system: if your brand is being discussed in communities in ways that conflict with your intended positioning, that signal will eventually surface in AI answers. Better to know early.

Monitoring community visibility is not the same as measuring positive sentiment. A domain can be highly cited because it hosts recurring debate about a category, not because it presents brands favorably. Source mapping and community intelligence work together to give a complete picture.

Experts are credibility infrastructure.

Professional network visibility matters differently across AI engines, but the direction is consistent: platforms that surface expert, practitioner, and B2B credibility contribute to citation visibility.

Building professional and expert presence means treating your organization's people as credibility assets, not just spokespeople. That includes maintaining accurate, detailed LinkedIn profiles for executives and subject-matter experts; developing visible expert bios that are indexed and findable; publishing research, commentary, and analysis that establishes genuine expertise on record; securing third-party profiles, awards, and listings that provide external validation; and ensuring that expert citations and references in credible publications point back to verifiable sources.

This workstream is particularly important for categories where trust is high stakes including professional services, healthcare, financial services, technology, and other regulated industries. These groups are ones where AI systems are more likely to draw on authoritative professional sources than on general consumer platforms. Executive and expert visibility in these categories is not a vanity metric. It is a credibility input into the source environments AI answers rely on.

Accuracy is a competitive advantage.

Reference and institutional sources, such as Wikipedia, government agencies, research institutions, academic publishers, remain a persistent layer in AI citation environments. This



is particularly true for factual, scientific, technical, and regulated categories. These sources are cited not because they are promotional, but because AI systems treat them as reliable anchors for factual claims.

The practical implication for brands is that accuracy and institutional credibility are not just compliance concerns. They are reputation infrastructure. A Wikipedia entry that contains outdated or inaccurate information about your organization is a citation risk. An absence from relevant institutional or regulatory databases is a credibility gap. A research citation that leads to a retracted paper or a broken link is an authority signal that works against you.

Reference and authority protection means auditing the factual layer around your brand regularly. The checklist for those activities includes checking key reference entries for accuracy, maintaining correct and current information across institutional and regulatory listings, developing credible research assets and white papers that can serve as citable sources, and building relationships with academic or institutional partners where appropriate to your category.

This workstream is not relevant at equal intensity for every brand or category. Consumer goods brands need it less than pharmaceutical companies. But for any organization operating in a category where AI systems are likely to cite factual, institutional, or regulated sources, protecting the accuracy and completeness of the reference layer is essential.

Visibility changes. Watch it regularly.

For the Semrush LLM *Citations Study,* the citation data spans 34 weekly reporting periods. Over that time, source sets shifted. Some domains gained significantly, others had dramatic declines while new sources entered the top tiers in specific markets and engines. These movements are not random. They reflect platform growth, model updates, changes in how AI providers weight different source types, and shifts in how communities discuss topics.

The implication is straightforward: ***AI citation visibility is not a static condition***. A brand that was well-represented in AI answers six months ago may not be today. A local source or competitor that was irrelevant in one period may have become a primary reference point in the next.

Recurring prompt-level monitoring should be set up by market, by engine, and by topic area relevant to your brand. Exercising this reporting rigor keeps citation intelligence current and actionable. This is not the same as tracking traditional media coverage or social mentions. It requires monitoring the specific prompts your audiences are likely to ask, across the engines they use, and observing which sources appear.



Treat sudden changes in citation patterns as a signal to investigate, not to immediately act on. A sharp movement in a domain's prompt share may reflect a genuine shift in how AI systems weight that source. It is important to exercise caution because it may reflect a temporary change in model behavior. The discipline is to notice, investigate, understand, and then respond with a measured strategy rather than reactive content production.

A Parting Word

The brands that will navigate this environment most effectively are not those that find clever ways to insert themselves into AI answers. A tactic that works today might not work tomorrow.

Just as was seen when traditional search matured, it weeded out and punished some of these behaviors. That same path will most likely be applied to AI engines so they can build and maintain trust with users that they are providing the most accurate and authentic results.

By creating supporting, genuine content through earned media, community conversation and other sources, those brands are the ones that build genuine credibility ecosystems. Thoughtful content generated in and by communities, expert networks, local markets, reference layers, and professional platforms will be what AI systems and human audiences find when they look for evidence.

That is not a new objective for communications. It is a more demanding version of the same one. The terrain has changed. The standard has not.



References

[1] There are several studies that have been completed about external v. owned sources including AirOps, asXiv, Omniscent, and Slate.

[2] Semrush LLM Citation Study: A deep analysis of citation-domain visibility across ChatGPT Search, Gemini, Google AI Mode, and Perplexity. This report analyzes domains cited directly in AI-generated answers. It replaces the previous clickstream-based proxy analysis and should not be compared as a continuation of the same metric.

[3] I/O 2026: Welcome to the agentic Gemini era

[4] Scaling AI for everyone

[5] Similarweb ChatGPT Traffic – snapshot taken in June 2026.