

AI Visibility Optimization: How Businesses Get Named by AI

J. Marcus Howard · President & CEO, Ethos Media Marketing & Consulting

Published: August 2, 2026 · Last updated: August 2, 2026

What is AI visibility optimization?

AI visibility optimization is the practice of making a business easy for artificial intelligence systems to find, verify, and name when people ask for recommendations. It has two parts: making the facts about a business consistent and machine-readable everywhere they appear, and building the independent outside references that AI systems treat as evidence the business is credible. The term covers, in one discipline, what generative engine optimization (GEO), answer engine optimization (AEO), and LLM SEO each describe in part.

Executive summary

For twenty years, being found online meant one thing: appearing in a list of links. That model is ending. A growing share of buying questions — who is the best dentist near me, which roofer should I trust, which med spa is worth the money — are now answered by artificial intelligence assistants that reply with a short written answer naming a handful of businesses. There is no page two. A business is either named or it is absent.

This paper explains, in plain language, how AI systems decide which businesses to name, what published research says actually moves that decision, and why the answer is infrastructure rather than a single trick. It draws only on sources that can be checked: a Gartner forecast, consumer research from Bain & Company, a peer-reviewed study from researchers at Princeton and their collaborators, and reporting from American Banker. Every claim traces to the original.

The key findings:

Finding 1. The question-and-answer shift is measured, not speculative. In February 2024, Gartner forecast that traditional search engine volume would drop 25% by 2026 as AI chatbots and virtual agents absorb queries. As of mid-2026 the full 25% has not materialized in total search volume, but the behavioral shift underneath it is documented and accelerating.

Finding 2. Consumers already trust the answer more than the links. Bain & Company's consumer research, published February 2025, found that 80% of consumers rely on AI-written results for at least 40% of their searches, that about 60% of searches now end without a click to any website, and that organic web traffic has fallen an estimated 15% to 25% as a result.

Finding 3. AI visibility is close to binary. A search results page shows ten or more options. An AI answer typically names a few. For a local business, the practical outcome is being recommended or being invisible — there is little middle ground.

Finding 4. Outside evidence is the strongest published lever. The first academic benchmark of this problem — by researchers at Princeton University, Georgia Tech, the Allen Institute for AI, and IIT Delhi, published at KDD 2024 — found that adding citations to sources, quotations, and statistics measurably increased how often content appeared in AI-generated answers, with visibility gains of roughly 30–40% and up to 115% for lower-ranked websites. Keyword stuffing, the classic search trick, did not work.

Finding 5. AI visibility has two requirements, and most businesses fail the second. A business must be machine-readable — its facts consistent, structured, and verifiable across the web — and machine-citable — referenced by independent third parties the AI treats as evidence. Most local businesses have partial readability and almost no citability.

Finding 6. Being named once is not a position. American Banker reported in February 2026 on testing

in which a company named in an AI assistant's first answer was gone from the conversation within a few follow-up questions. AI visibility must be maintained the way a reputation is maintained, not won once like a ranking.

Finding 7. Single tactics fail because the inputs are spread across the whole footprint. The signals AI systems weigh — consistency of facts, reviews, structured data, independent coverage, recency — live in different places. No one fix touches enough of them. The durable response is infrastructure that keeps all of them current at once.

1. The shift: from ten links to one answer

Search built the modern local economy on a simple contract: a person types a question, the engine returns a ranked list, and businesses compete for position on that list. Everything the marketing industry calls search engine optimization — SEO — exists to win that competition.

The contract is being rewritten. Large language models — the technology behind assistants like ChatGPT, Gemini, Claude, and Perplexity — answer questions directly. Ask one of them for the best pool builder in a county and it does not return ten links. It returns a paragraph, and the paragraph names names.

Gartner saw this coming before most of the market did. In a February 2024 forecast, the research firm predicted that by 2026, traditional search engine volume would drop 25%, with search marketing losing share to AI chatbots and other virtual agents [1]. Gartner's analyst framing was blunt: generative AI tools are becoming substitute answer engines, replacing queries people used to type into a search box.

It is worth being honest about where that forecast stands, because honesty is the standard this paper holds itself to. As of mid-2026, total search volume has not fallen by the full 25%; people still search, and the largest search engine has folded AI answers into its own results rather than losing the queries outright. But the direction is not in dispute, and the behavioral data underneath it is stark.

In February 2025, Bain & Company published consumer research on exactly this behavior. The findings: 80% of consumers now rely on AI-written results for at least 40% of their searches, and about 60% of searches end without the user clicking through to any website at all. Bain estimates organic web traffic has fallen 15% to 25% as a direct consequence [2].

Read those numbers together and the picture is simple. The question still gets asked. The answer still gets given. What is disappearing is the visit — the moment when a person lands on a website and the business gets a chance to make its own case. Increasingly, the case is made for you, by a machine, in a sentence or two, before anyone sees your site.

For a high-ticket local business — a dental practice, a remodeler, a marine dealer, a law firm — that sentence is the new front door. AI visibility optimization is the discipline of making sure your business is in it.

The economics deserve one more sentence, because they explain who should care most. For a business selling a forty-dollar item, one recommendation is a transaction. For a business selling a full-mouth restoration, a kitchen, a roof, or a retainer, one recommendation can be five or six figures — and the person asking an assistant that question is usually closer to a decision than a person scrolling links. The higher the average ticket, the more each appearance in the answer is worth, and the more expensive each absence quietly becomes.

2. How AI systems decide who to name

To influence the answer, you have to understand how it gets written. In plain terms, an AI assistant draws on two pools of information when it answers a question about businesses.

The first is what the model learned during training. Large language models are built by processing enormous amounts of public text — websites, articles, directories, reviews, published research. Businesses that appear often, consistently, and in credible places during that process are the ones the model "knows." A

business with a thin, inconsistent, or contradictory public record barely registers.

The second is what the system looks up in the moment. Most modern assistants also retrieve live information when asked — pulling from search indexes, listings, review platforms, and news — and then compose an answer from what they find. Here the system behaves like a fast, skeptical researcher: it favors sources it can verify, facts that agree with each other, and businesses that independent third parties have written about.

It helps to be concrete about where the looking happens. When a person asks an assistant about local businesses, the systems commonly draw on business listings and maps data, review platforms, the business's own website, local and trade press, directories, and whatever published material mentions the business by name. Each of those surfaces is a place the record can agree with itself — or contradict itself. An old address on one directory, a defunct phone number on another, a service list that differs between the site and the listing: to a machine assembling an answer under uncertainty, every contradiction is a reason to name someone else.

Both pools reward the same underlying qualities. Consistency: the business's name, address, phone number, services, and story match everywhere they appear. Structure: information is presented in formats machines parse cleanly, including structured data — the labeled code on a webpage that tells a machine "this is the business name, this is the address, these are the hours." Verification: claims are confirmed by outside sources, not just asserted by the business itself. And recency: the record is current, because systems discount stale information.

One more concept earns its place in plain language: the entity. Modern systems do not primarily understand a business as a website; they understand it as an entity — a distinct thing with attributes: a name, a location, services, people, a history, and connections to other entities. Everything published about the business either sharpens that entity or blurs it. A distinctive, consistently used name sharpens it. Three spellings of the company name, a founder credited under different titles, or a brand that collides with an unrelated company's name all blur it — and a blurred entity is one a cautious system declines to recommend. Entity clarity is unglamorous, foundational work, and it is upstream of everything else in this paper.

What makes this environment genuinely different from search rankings is how unstable position is. American Banker — the trade publication covering the U.S. banking industry — reported in February 2026 on testing of what generative AI tools say about financial institutions. One finding stands out for every business owner: a company named in the AI's first answer was gone from the conversation within a few follow-up questions, displaced as the user refined what they were asking [4].

Rankings move in weeks. AI answers can reshuffle inside a single conversation. That volatility is not a reason to ignore the channel — it is the reason visibility has to be built on signals that persist across every version of the question: verified facts and independent references, not a lucky mention.

3. What the research says actually works

Most advice about AI visibility is opinion. There is, however, a published, peer-reviewed starting point, and its results are specific enough to act on.

In late 2023, researchers from Princeton University, Georgia Tech, the Allen Institute for AI, and IIT Delhi published the first systematic study of how content choices affect visibility in AI-generated answers. The paper, titled "GEO: Generative Engine Optimization," was peer-reviewed and presented at KDD 2024, one of the leading data science conferences [3]. The team built a benchmark of 10,000 queries and tested which changes to a webpage made its content more likely to be used — and credited — in an AI engine's answer.

Three changes stood out, and none of them are exotic:

Citing sources. Adding references to credible outside sources produced the largest gains — in the study's tests, visibility improvements of roughly 30–40% on the study's measures, and up to 115% for websites that ranked lower in traditional search. Read that carefully: the sites that benefited most from verifiable outside evidence were the ones that were not already winning. That is the single most hopeful published finding for

small businesses competing against bigger names.

Adding quotations. Direct quotations from relevant sources — experts, named people, published material — measurably increased the chance content was drawn into answers.

Adding statistics. Concrete numbers from checkable sources outperformed general claims. Machines, like careful readers, prefer statements that can be verified.

Just as important is what failed. Keyword stuffing — repeating the target phrase over and over, the oldest trick in search — produced no meaningful gain in the study [3]. The systems writing AI answers are built to detect substance, and they discount its absence.

The pattern across all of it is one idea: evidence. AI answers are assembled the way a cautious researcher writes — preferring material that cites, quotes, and quantifies. A business that wants to be named needs to look, to a machine, like the best-evidenced answer to the question.

For a local business, translating the findings is straightforward. Service pages and guides should cite checkable outside sources — industry bodies, published standards, government data — rather than asserting alone. Pages should carry real quotations from named people: the owner, the lead technician, a satisfied client who agreed to be quoted. Claims should be numeric where a checkable number exists — years in operation, licenses held, jobs completed where documentable — and general where it does not, because an unverifiable number is worse than none. None of this is exotic writing. It is the difference between a page that asserts and a page that demonstrates, and the research says the machines can tell them apart [3].

4. The two requirements: readable and citable

Put the mechanics of Section 2 next to the research of Section 3 and the whole discipline reduces to two requirements.

Requirement one: be machine-readable. Every fact about the business — name, address, phone, services, hours, pricing posture, the story of who you serve — must be consistent and current everywhere it appears: the website, business listings, directories, review platforms, maps. Structured data should label those facts so machines parse them without guessing. Reviews should be active and recent, because they are among the few third-party signals nearly every retrieval system checks. Readability is the floor. It determines whether an AI system can confidently state who you are.

Requirement two: be machine-citable. Independent third parties — publications, industry outlets, local media, research — must reference the business, because outside references are what these systems treat as evidence. This is the requirement the Princeton-led study measured most directly, and it is the one most local businesses fail completely. A business can have a flawless website and perfect listings and still never be named, because nothing independent of the business confirms it matters.

In practice, readability is a finite checklist. One business name, one address, one phone number, everywhere — including the old listings nobody remembers creating. Structured data on the site marking the organization, its location, services, and hours. A review flow that is active this month, not last year, with responses that show the business is present. Service pages that state, in plain sentences, what is done, where, and for whom. Hours, service area, and pricing posture current. A machine reading all of it should reach the same set of facts from every direction.

Citability is not a checklist; it is an accumulation. It includes coverage in local and trade outlets, published material under the business's expertise — guides, research, data worth referencing — commentary contributed where the industry's questions are being discussed, and distribution that places the business's knowledge on platforms and publications it does not own. What unites every form is the one property the business cannot manufacture alone: someone else said it. That is why citability is the scarcer asset, the slower one to build, and — per the research — the one that moves the answer most [3].

The two requirements are not interchangeable, and they fail differently. A business that is citable but not readable confuses the machine: coverage exists, but the facts do not line up, so the system hedges or names

a competitor whose record is cleaner. A business that is readable but not citable is legible and ignorable: the machine knows exactly who you are and has no evidence you are the answer. Only both together produce the outcome that matters — being named, and staying named as the questions keep coming.

Everything sold under the labels GEO, AEO, and LLM SEO is a partial view of these two requirements. Naming them plainly — readable and citable — makes the work assignable, measurable, and honest.

5. Why single tactics fail

The market's instinct with every new channel is to look for the one fix: a plugin, a listing service, a burst of blog posts, one press release. In this channel the instinct fails for a structural reason: the signals AI systems weigh do not live in one place.

Consistency lives across dozens of listings and pages. Structure lives in website code. Reviews live on third-party platforms. Citations live on other people's websites and in the press. Recency lives in all of them at once, and decays in all of them at once. A tactic touches one surface. The decision is made across all of them.

The pattern shows up in the three purchases owners most often try. A one-time press release creates a reference, then ages; six months later it is one stale data point in a record that has otherwise not moved. A listings-sync subscription keeps the facts tidy — genuine readability value — while doing nothing about the absence of independent evidence, which is usually the binding constraint. A burst of AI-written blog posts adds volume to the one surface the business already controls, which is precisely the surface the research says machines discount when nothing outside confirms it [3]. Each purchase is rational in isolation. Each fails alone, for the same reason: it improves one signal while the decision is made across all of them.

This is also why the work cannot be finished. The American Banker testing described above showed position dissolving inside a single conversation [4]; the same logic applies over months, as fresher, better-evidenced competitors enter the record. A business that builds readability once and stops will drift out of date. A business that earns one round of coverage and stops will watch its evidence go stale.

The correct mental model is not a campaign. It is infrastructure — the same way payroll, insurance, and bookkeeping are infrastructure. Systems that keep the facts consistent, the structure current, the reviews flowing, and the independent references accumulating, on a schedule, whether or not anyone is thinking about them that week.

6. A framework view: where this work lives in a growth system

This section states our perspective as practitioners; the evidence above stands on its own without it.

Our firm organizes client growth through the FOCAS Growth Framework™ (Foundation, Operations, Conversion, Amplification, Scale), and delivers it for local businesses as the FOCAS Autonomous Marketing System (FOCAS AMS™) — AI-led, human-supervised, and running all five stages as one managed service. Mapped against the two requirements, AI visibility is built in the first stage and won in the fourth.

Foundation is where a business becomes machine-readable. Consistent business facts across every listing and page, structured data that labels them, active review flow, a website whose claims agree with the public record. When Foundation work is done properly, an AI system can verify who the business is, where it operates, and what it does — without ambiguity.

Amplification is where a business becomes machine-citable. Earned coverage in independent outlets, published material worth referencing, multi-format distribution that puts the business's expertise in places machines index and trust. This is deliberate, ongoing work — and it is precisely the lever the published research identifies as strongest, especially for businesses that are not already dominant [3].

The middle stages matter here for a reason owners feel immediately. Operations is the systems layer — the follow-up, scheduling, and response infrastructure that answers when opportunity arrives. Conversion is the discipline of turning attention into booked, paying work. AI visibility raises the stakes on both: when an

assistant recommends a business, the person arriving is further along, more decided, and less patient than a browser clicking a link. A missed call or a two-day reply squanders a recommendation that was hard to earn. Winning the answer and fumbling the arrival is the most expensive failure in the sequence.

Scale compounds what works. But for visibility specifically, the failure we see most often has a name: the Authority Ceiling, a term we introduced in the book of the same name [5]. It is the point at which a business's own effort stops moving results — the website is good, the content is real, the service is excellent — because everything the business can say about itself has been said, and nothing independent has been built to confirm it. Below the ceiling, more self-published effort produces less and less. The ceiling breaks from the outside: independent references, accumulated deliberately. In the AI era, that is not a branding preference. It is, per the research above, the measured mechanism by which answer engines decide whom to trust.

From the owner's seat, the ceiling has a familiar feel. The website gets rebuilt and nothing moves. The blog gets consistent and nothing moves. Rankings plateau on the terms that matter, referrals quietly carry the business, and every additional hour of self-published effort returns less than the hour before it. Owners usually diagnose this as a content problem or a budget problem. It is neither. It is an evidence problem: the business has said everything it can say about itself, and no one independent has said anything at all.

Inside that framework, our practice tests a business against three layers of authority, in order. **Entity Recognition:** can AI systems tell, without ambiguity, who the business is, where it operates, and what it does? **Authority Signals:** has anyone independent — publications, institutions, other experts — confirmed any of it? **Expertise Alignment:** does that evidence match the specific things customers actually ask about? A business can pass the first layer and still go unnamed; most established local businesses fail at the second.

The build unit we use against the second and third layers is the Amplified Media Placement (AMP) — a single piece of owner expertise produced in seven or more formats and distributed to a network of more than 400 independent publications, so that one insight becomes many independent references. AMP campaigns run on a schedule rather than in bursts, typically opening with a 90-Day Authority Ceiling Breakthrough Sprint, because the working measure in this channel is authority velocity: how quickly a business accumulates and maintains independent evidence relative to its competitors. The names are ours; the mechanism they operationalize is the one the research in Section 3 measured.

7. What this looks like in practice: a client example

Elite Professional Dog Training is a Central Florida dog training company whose owner has trained dogs professionally since 1990 and has operated the business since 2007 — an established, well-run local business of exactly the kind this paper describes: real expertise, real results, and a public record that undersold both.

The work followed the two requirements in order. First, readability: the business's facts, listings, and site were brought into consistent, structured order. Then, citability: the owner's expertise was developed into published material and distributed through independent outlets in AMP campaigns, building the outside record that had never existed.

In a client testimonial published on our framework site, the owner describes the result in his own terms: coverage syndicated across more than 400 publications, growth strong enough that his banker asked what had changed — and, during the engagement, expansion from one location to a second. Decades of skill did not change. What changed is that the public record finally carried evidence of it — the kind machines and people both count.

The profile is worth naming because it is the standard one. An established business, years of genuine skill, a decent website, steady referrals — and almost nothing independent in the public record. That combination describes the majority of high-ticket local businesses in America. It is also, per Section 3, precisely the profile the research says has the most to gain from evidence-building: the largest measured improvements went to the players who were not already winning [3].

One client example proves possibility, not statistics, and we present it as exactly that. The measured case for the mechanism is the published research in Section 3. The example shows what the mechanism looks like when a real local business runs it.

8. How to check where you stand

A business owner can get a rough read on both requirements in an afternoon, without tools.

Ask the machines. Open two or three AI assistants and ask the questions your customers ask: best [your service] in [your area], who should I trust for [the job], is [your business name] reputable. Note who gets named, whether you appear, and what the assistant says when asked about you directly. Then ask two follow-up questions and watch whether the names change — remembering that position can dissolve within a conversation.

Audit your readability. Search your own business name, address, and phone number. Do they match on every listing, or are there old addresses and dead numbers in the record? Does your site state plainly what you do and where? Are your reviews recent, or does the trail stop months ago?

Audit your citability. Search your business name in quotation marks minus your own website. What is left — what exists about you that you did not publish yourself? For most local businesses the honest answer is: almost nothing. That gap is the Authority Ceiling, located.

If the read sends you shopping for help, five questions separate providers quickly. First: who does the work — do you execute the fixes and build the references, or do you measure, score, and hand me a report? A growing class of providers sells measurement and advice; measurement locates the gap, and only execution closes it, so ask which one you are buying. Second: what exactly will exist in ninety days that does not exist today — listings corrected, structured data live, references published where? Third: how is progress measured against my actual competitors, not against a generic score? Fourth: what happens when the work stops — does anything I paid for persist? Fifth: show me a business like mine, before and after. A provider with good answers to all five is selling infrastructure. A provider with good answers to none of them is selling a dashboard.

Our practice scores that read as the Authority Ceiling Index (ACI), built from three component scores: a Technical Readiness Score (TRS) — can AI systems understand you; a Content Authority Score (CAS) — have you answered the questions buyers actually ask; and an Authority Amplification Score (AAS) — does anyone else vouch for you. The TRS is this paper's readability requirement in score form; the CAS and AAS together are its citability requirement, split into the record you publish and the record others confirm.

For a systematic version of that read, our firm operates a free diagnostic through Catalyst Pro that benchmarks a business's authority signals — including AI citation rate — against its top three competitors, in about sixty seconds, with no credit card: trycatalystpro.com/diagnostic. It tells you which of the two requirements is holding you under the ceiling, and by how much.

9. What comes next

Two developments make this work more valuable over time, not less.

The answers are becoming the interface. AI answers are moving inside the tools people already use — phones, browsers, cars, voice assistants — and each move shrinks the moment where a list of options appears. The fewer options shown, the more it matters to be the name.

The assistants are starting to act. The industry is building agents — AI systems that do not just recommend a business but book the appointment, request the quote, or place the call. When the assistant completes the transaction, the recommendation is the sale. The record that gets a business named today is the same record that gets it chosen by an agent tomorrow.

One last point of orientation. The answers already exist. Ask an assistant about your category in your market tonight and it will reply — naming someone, drawing on whatever record it can find, including whatever

exists about you. Declining to participate does not remove a business from the conversation; it only guarantees the business has no influence over its part in it. Silence, in this channel, is also an answer — usually a competitor's name where yours could be.

Both trends compound the same asset: a verified, independently evidenced public record. That asset appreciates. Businesses that begin building it now are accumulating exactly what the next interface will read.

Questions and answers

What is AI visibility optimization?

AI visibility optimization is the practice of making a business easy for AI systems to find, verify, and name when people ask for recommendations. It combines making the business's facts consistent and machine-readable everywhere they appear with building the independent outside references AI systems treat as evidence of credibility.

How is AI visibility optimization different from SEO?

SEO competes for position in a list of links; AI visibility optimization competes to be named in a written answer. They overlap — clean structure and credible content help both — but AI answers weigh verification and independent references more heavily, and they name only a few businesses instead of ranking many.

Why does my business rank on Google but not show up in ChatGPT?

Ranking reflects how well your site competes in a search index. Being named in an AI answer depends on whether the system can verify your facts and find independent references confirming your credibility. A business can rank well on its own content and still have almost no outside evidence — which is the most common reason it goes unnamed.

How do I get my business mentioned by ChatGPT and other AI assistants?

Meet two requirements. Make the business machine-readable: consistent name, address, phone, and services across every listing, structured data on your site, and active reviews. Then make it machine-citable: earn references from independent publications and outlets. Published research found that verifiable outside evidence — citations, quotations, statistics — measurably increases visibility in AI answers, with the largest gains going to businesses that are not already dominant.

How long does it take to build AI visibility?

Readability work — fixing facts, listings, and structure — can be completed in weeks and is often reflected quickly. Citability compounds over months, because independent references accumulate with each round of coverage. Both require maintenance: AI systems discount stale records, and position can shift as competitors build their own evidence.

Can I just buy a tool that tracks my AI visibility?

Tracking tools answer one question: where do I stand? They do not correct a listing, publish structured data, earn a reference, or keep any of it current. Measurement is the map; execution is the trip. For most local businesses the constraint is not knowing the gap exists — it is having the infrastructure that closes it and keeps it closed.

Does AI visibility matter for a small local business, or only for big brands?

The published research points the other way: the largest measured gains from adding verifiable outside evidence went to websites that ranked lower in traditional search — improvements of up to 115% in the study's tests. Big brands already have citations; that is much of why they are big. For a small business, evidence-building is the rare channel where the underdog's next unit of effort moves the result more than the leader's.

What should I do first?

Sequence matters. First, fix readability — reconcile every listing, correct the facts, add structured data, restart the review flow — because citations pointing at a contradictory record are wasted. Second, begin building citability on a schedule: published expertise, earned references, distribution beyond your own site. Third, put maintenance on a calendar, because both requirements decay. Businesses that reverse the order pay for coverage that lands on a foundation machines cannot verify.

How do I measure AI visibility?

Ask the major assistants your customers' actual questions on a schedule and log who gets named. Track your citation footprint — the independent references to your business that you did not publish yourself. Benchmark both against your closest competitors, because visibility is relative: the machine is choosing among you.

What is the Authority Ceiling?

The Authority Ceiling is the threshold of independent, verifiable evidence below which AI systems cannot reliably identify, verify, or recommend a business. Below it, more self-published effort — a better website, more blog posts — returns less and less, because the missing ingredient is evidence from sources the business does not control. The term comes from the book *The Authority Ceiling* and names the barrier most established local businesses are under without knowing it.

What is the Authority Ceiling Index (ACI)?

The Authority Ceiling Index (ACI) is a scoring instrument, introduced in the book *The Authority Ceiling*, that measures how far a business sits above or below the ceiling. It is built from three component scores: a Technical Readiness Score (TRS), which measures whether AI systems can understand the business; a Content Authority Score (CAS), which measures whether the business has answered the questions buyers actually ask; and an Authority Amplification Score (AAS), which measures whether anyone independent vouches for it. Read together, the three scores show which requirement — readability or citability — is holding a business back.

What is an Amplified Media Placement (AMP)?

An Amplified Media Placement (AMP) is a content campaign format developed by Ethos Media Marketing & Consulting in which one piece of a business owner's expertise is produced in seven or more formats — article, video, audio, visual, and more — and distributed to a network of more than 400 independent publications. Its purpose is to convert genuine expertise into the independent, citable references that AI systems count as evidence.

A note on method

Every statistic in this paper traces to its original source, listed below; nothing is cited secondhand. Where a forecast has not fully materialized, the paper says so. A companion methodology paper — covering how we benchmark AI visibility for local businesses, including query design and measurement cadence — is forthcoming from our research program and will be deposited with a permanent digital object identifier (DOI).

References

1. Gartner, Inc. "Gartner Predicts Search Engine Volume Will Drop 25% by 2026, Due to AI Chatbots and Other Virtual Agents." Press release, February 19, 2024. [gartner.com](https://www.gartner.com)
2. Bain & Company. "Consumer reliance on AI search results signals new era of marketing." Press release, February 19, 2025; companion report "Goodbye Clicks, Hello AI: Zero-Click Search Redefines Marketing." [bain.com](https://www.bain.com)
3. Aggarwal, P., et al. "GEO: Generative Engine Optimization." Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD 2024). arxiv.org/abs/2311.09735
4. American Banker. "Do you know what gen AI is saying about your bank?" February 2026. [americanbanker.com](https://www.americanbanker.com)

5. Howard, J. M. *The Authority Ceiling: Breaking Through the Invisible Barrier That Keeps Your Business Unknown — A Complete Strategic Guide to AI Visibility, Digital Authority, and Category Domination*. Catalyst Pro / Ethos Media Marketing & Consulting LLC, first edition, 2025. trycatalystpro.com

How to cite this paper

Howard, J. M. (2026). AI Visibility Optimization: How Businesses Get Named by AI. Ethos Media Marketing & Consulting. <https://ethosm2.com/ai-visibility-optimization/>

Go deeper

Related reading from our Learning Center:

- [More on AI visibility and local marketing](#)

About the author

J. Marcus Howard is President & CEO of Ethos Media Marketing & Consulting LLC, the architect of the FOCAS Growth Framework™, and the author of *The Authority Ceiling* [5]. He is a repeat presenter on the BridgeTECH track at the Bridge Conference (2024, 2026); his 2026 session, "Breaking the Authority Ceiling: How Small Nonprofits Use AI-Powered Content Distribution to Compete with Major Organizations," appears on the official conference agenda ([2026 agenda](#)).

About Ethos Media Marketing & Consulting

Ethos Media Marketing & Consulting LLC (EthosM2) builds digital marketing infrastructure for brick-and-mortar local businesses, organized through the FOCAS Growth Framework™ (Foundation, Operations, Conversion, Amplification, Scale). The firm's research and product releases are announced via national newswire distribution. Learn more at ethosm2.com and stayfocas.com.

Close the Digital Authority Gap

You now know how the answer gets written. The remaining question is where your business stands in it — which of the two requirements you already meet, and which one is holding you under the ceiling.

The Authority Diagnostic answers that in about sixty seconds, free, with no credit card: it benchmarks your authority signals and AI citation rate against your top three competitors and shows you exactly where the gap is.

[Run the free Authority Diagnostic](#)