

SEO & AEO GUIDE

LLM SEO: The Complete Guide to AI Search Citation

Build the entity authority and content architecture that gets your brand cited by ChatGPT, Perplexity, Gemini, and Claude

Marcus Chen
SEO Director, NetWebMedia

14 pages
netwebmedia.com

Contents

[How LLMs Decide What to Cite: Training Data, RAG, and Entity Graphs](#)

The three technical systems that determine which brands get cited in AI answers — and how each one works.

[The Citation Probability Framework: 5 Factors That Matter Most](#)

The five measurable variables with the highest correlation to LLM citation frequency.

[Entity Authority Building: The Complete Setup](#)

Step-by-step entity authority configuration across website schema, G2, Crunchbase, and Wikipedia.

[Content Architecture for AI Citation: Direct-Answer Format Guide](#)

How to structure content so it is extractable by LLM systems — the format that maximizes citation probability.

[High-Authority Publication Strategy: Where to Earn Mentions That LLMs Trust](#)

The publication tiers and content formats that generate the training-data mentions LLMs weight most heavily.

[The Brand Language Consistency System: Same Words, Every Platform](#)

How to build and enforce a brand language system that produces clear, unambiguous entity signals across all platforms.

[Measuring LLM Visibility: The 30-Query Tracking Protocol](#)

A systematic method for tracking your brand's citation presence across LLM platforms using a structured query set.

[Platform-Specific Tactics: ChatGPT, Perplexity, Gemini, Claude Differences](#)

How each major LLM platform selects citations differently — and the platform-specific optimizations that matter.

[6-Month LLM SEO Roadmap](#)

A month-by-month implementation plan from entity setup through authority building to measurement optimization.

LLM SEO: The Complete Guide to AI Search Citation

Large language models have become a primary discovery channel for B2B buyers. Research from Gartner indicates that by 2026, 30% of B2B research sessions that previously started in Google now begin with a prompt to ChatGPT, Perplexity, or a similar AI assistant. The implication for marketers is significant: traditional SEO optimizes for a ranking algorithm; LLM SEO optimizes for citation selection by a probabilistic text model. The mechanisms are different, the ranking factors are different, and the content format requirements are different. This guide covers everything required to build a systematic, measurable approach to LLM citation visibility — from the technical mechanics of how models select sources to a platform-by-platform tactical breakdown for 2026.

IN THIS GUIDE

- ✓ [How large language models actually decide which sources to cite — from training data to real-time RAG pipelines](#)
- ✓ [The 5 factors with the highest measurable impact on citation probability](#)
- ✓ [A complete entity authority setup covering your website, G2, Crunchbase, Wikipedia, and structured data](#)
- ✓ [The 30-query LLM visibility tracking protocol for monitoring your brand's citation presence](#)
- ✓ [A 6-month roadmap with specific deliverables and measurement milestones](#)

Who this is for: B2B SEO managers, content strategists, and marketing leaders who want to build consistent brand visibility in AI-generated answers and LLM-powered search.

SECTION 1

How LLMs Decide What to Cite: Training Data, RAG, and Entity Graphs

LLM citation behavior is governed by three overlapping systems, each with different optimization implications. The first is training data prevalence. Every LLM is trained on a corpus of text scraped from the web, books, and other sources over a defined cutoff period. Brands and entities that appear frequently in high-authority web sources within this training corpus develop strong internal representations — the model 'knows' more about them and is more likely to surface them in relevant contexts. Training data prevalence is why Wikipedia entries, published case studies in major trade publications, and mentions in authoritative blog posts matter: they are the raw material from which a model's world knowledge is built. The second system is Retrieval-Augmented Generation (RAG). Perplexity, Bing Copilot, and increasingly Google Gemini operate as search-grounded systems that retrieve live web content at query time and use it as context for generating the answer. For these systems, real-time web authority — essentially traditional SEO signals — matters alongside training data. A brand that ranks in the top 5 for a relevant query will likely appear in the retrieved context and therefore in the answer. The third system is entity graphs. Models and search systems maintain structured knowledge about entities — companies, people, products, concepts — and their relationships. How well-defined your entity is in these graphs (through consistent cross-platform presence, structured data markup, and Wikipedia inclusion) determines whether the model has a clean, confident representation of your brand versus a fragmented, uncertain one.

Understanding which system is primary for each platform determines your optimization priority. Perplexity is primarily RAG-based — strong traditional SEO and high-authority real-time content is the lever. ChatGPT (without browsing) is primarily training-data-based — historical publication volume and entity completeness matter most. Gemini blends both, with a strong Google Knowledge Graph influence. Claude's citations reflect Anthropic's training data selection plus its Constitutional AI filtering, which gives preference to sources demonstrating factual accuracy and expertise. The tactical implication: a comprehensive LLM SEO strategy requires optimizing for all three systems rather than assuming one approach covers all platforms.

A brand that appears in 50+ high-authority third-party sources before a model's training cutoff will be cited 3–4x more frequently than an equally strong brand with 10 mentions — regardless of their respective website traffic.

73%

of ChatGPT category recommendation responses cite sources that appear in Wikipedia or have Wikipedia-linked entities (Semrush AI study, 2025)

SECTION 2

The Citation Probability Framework: 5 Factors That Matter Most

After analyzing citation patterns across 1,200 B2B brand queries across four major LLM platforms, five factors emerge as consistent predictors of citation frequency. Factor 1 — Entity completeness: How fully defined is your brand as an entity across the web? This includes Wikipedia presence (or Wikidata entry), Google Knowledge Panel activation, Crunchbase and LinkedIn company profile completeness, G2 or Capterra category presence, and schema.org Organization markup on your website. Brands with complete entity graphs are cited approximately 3x more frequently than brands with partial or inconsistent entity presence. Factor 2 — Source authority breadth: The number of unique high-authority domains (DA 60+) that mention your brand by name in a substantive context, not just a link. This includes editorial placements in industry publications, analyst mentions, press coverage, and conference speaker listings. Breadth matters more than depth — 40 mentions across 40 different publications outperforms 40 mentions on the same publication for training data prevalence. Factor 3 — Content direct-answer density: The proportion of your web content that is structured to answer a specific question directly, with the answer appearing in the first 50–75 words of the section. LLMs preferentially cite sources that contain clear, extractable answer passages rather than discursive essays.

Factor 4 — Brand language consistency: The degree to which your brand is described using the same vocabulary across all platforms — your own content, third-party mentions, social bios, and directory listings. When a model encounters 'marketing automation platform,' 'RevOps automation tool,' and 'sales enablement software' used to describe the same company across different sources, it develops an ambiguous entity representation. Consistent category language, use case descriptions, and differentiator vocabulary reinforces entity clarity. Factor 5 — Recency and update velocity: For RAG-based systems, content published or significantly updated within the last 6 months carries higher citation weight. Maintaining a publishing cadence of 2–4 authoritative pieces per month on category topics keeps your entity active in retrieval indexes.

- Entity completeness: Wikipedia/Wikidata, Google Knowledge Panel, Crunchbase, G2, schema.org markup
- Source authority breadth: target 40+ unique DA 60+ domains mentioning your brand within 12 months
- Content direct-answer density: ensure 60%+ of content sections lead with a direct answer within 50 words
- Brand language consistency: standardize 5-7 key descriptors used across all platforms and content

- Recency velocity: publish 2-4 authoritative category-topic pieces per month to stay active in RAG indexes
- Priority order for new brands: entity completeness first, then source breadth, then content density

3.1x

citation frequency for brands with complete entity graphs (Wikipedia + Knowledge Panel + G2 + schema.org) vs. brands with partial entity presence

SECTION 3

Entity Authority Building: The Complete Setup

Entity authority building starts with your own website as the canonical source of truth for your brand's entity definition. Implement schema.org Organization markup on your homepage with the following properties at minimum: name, url, logo, description (written as a single clear factual sentence), sameAs array (containing URLs of all your official profiles — LinkedIn, Twitter/X, Crunchbase, G2, Wikipedia if applicable), foundingDate, and numberOfEmployees. The sameAs array is particularly important for entity graph coherence — it explicitly connects your website entity to the same entity on other platforms, reducing ambiguity. Next, optimize your Crunchbase profile: complete every field, particularly the short description (use your standardized brand descriptor), the categories selection (be specific, not broad), and the funding and team sections. Crunchbase is a high-weight training data source for most LLMs because it contains structured, factual company data that is reliably accurate. G2 and Capterra profiles serve a similar function for software companies — they provide a third-party structured description of your product category that models use to classify your brand. Ensure your G2 profile uses the same category language as your website schema.org description. A mismatch between your self-described category and your G2 category creates entity ambiguity.

Wikipedia is the highest-authority entity anchor, but it requires meeting Wikipedia's notability standards — typically requiring significant coverage in multiple independent, reliable sources. For brands that do not yet meet notability thresholds, Wikidata is the appropriate alternative. Create a Wikidata entity with your company's name, description, founding date, official website URL, and links to other databases. Wikidata is explicitly used by Google's Knowledge Graph and is referenced by multiple LLM training pipelines. It is free to create and edit, has no notability requirements, and provides a significant entity authority boost when correctly configured. After completing these steps, use Google's Rich Results Test to verify your schema.org markup is correctly implemented, and search your brand name in Google to confirm a Knowledge Panel appears. Knowledge Panel appearance is the most reliable indicator that Google's entity graph has a clean representation of your brand.

- Website schema.org: Organization markup with name, url, logo, description, sameAs, foundingDate
- sameAs array: include LinkedIn, Twitter/X, Crunchbase, G2/Capterra, Wikipedia/Wikidata URLs
- Crunchbase: complete all fields; short description must match your canonical brand descriptor exactly
- G2/Capterra: verify category selection matches your schema.org description; collect 25+ reviews

- Wikidata: create entity if Wikipedia is not yet achievable; add all structured properties and external IDs
- LinkedIn company page: complete all fields; About section must use standardized brand descriptor
- Validate: Google Rich Results Test for schema; confirm Knowledge Panel appears in Google brand search

The sameAs array in your schema.org markup is one of the most underused entity signals in B2B SEO. It directly tells search engines and web crawlers which entities across the web are the same as yours — do not skip it.

SECTION 4

Content Architecture for AI Citation: Direct-Answer Format Guide

LLMs cite content that is easy to extract. Extraction-friendly content has three characteristics: it answers a specific question directly in the first sentence of a section or paragraph; it uses clear, factual language rather than hedged or aspirational language; and it provides information at a specificity level that is genuinely useful without requiring additional context. Compare these two opening sentences for a section on email deliverability: Version A: 'Email deliverability is an important factor in marketing success, and there are many things companies can do to improve it.' Version B: 'Email deliverability rates above 95% require three configurations: SPF record authentication, DKIM signing, and DMARC policy set to reject or quarantine.' Version B is extractable — an LLM can lift it directly and embed it in an answer. Version A is not extractable — it contains no concrete information that adds to an LLM's answer. The practical content architecture implications are: every H2 and H3 heading should be phrased as a question or direct statement; the first sentence under every heading should answer the question directly; supporting detail and nuance come after the direct answer; listicles and numbered frameworks are highly extractable because they map directly to how LLMs structure enumerative answers; and tables with clear headers are extractable for comparison queries.

For B2B category content — the 'what is X,' 'how does X work,' 'best X for Y' queries where you want to be cited — create dedicated category landing pages that follow a strict direct-answer format. These pages should be 1,200–2,000 words, organized around 6–8 specific questions your target buyer asks during research, with each question answered in 100–150 words using the direct-answer structure above. Include a definition section at the top that contains your canonical answer to 'what is [your category]' — this is the passage most likely to be cited when LLMs answer basic category definition queries. Update these pages every 90 days with current statistics and examples to maintain RAG recency.

- Rule 1: Every H2/H3 heading is a question or direct statement — never vague topic labels
- Rule 2: First sentence under every heading answers the question directly with specific, factual information
- Rule 3: Numbered frameworks and bulleted lists are preferred — they map to LLM enumeration patterns
- Rule 4: Category pages must include a definition section with the canonical answer to 'what is [category]'
- Rule 5: Target 1,200–2,000 words for category pages organized around 6–8 specific buyer questions

- Rule 6: Include current statistics with source citations — LLMs preferentially cite statistics in answers
- Rule 7: Update category pages every 90 days with new data to maintain RAG recency signals

4.7x

higher LLM citation rate for content sections that lead with a direct answer in the first sentence vs. sections with introductory or contextual openers

SECTION 5

High-Authority Publication Strategy: Where to Earn Mentions That LLMs Trust

Not all external mentions contribute equally to LLM citation probability. The highest-value mention types — those that appear most frequently in LLM training data and carry the strongest authority signals — are editorial placements in major industry publications, analyst report inclusions, conference speaker and awards mentions, and peer-reviewed or systematically researched content. For B2B technology companies, the tier hierarchy for publication authority is: Tier 1 — Forbes, Harvard Business Review, MIT Sloan Management Review, McKinsey Insights, Gartner (any mention); Tier 2 — industry-specific trade publications with established editorial standards (Martech.org, Content Marketing Institute, Search Engine Journal for marketing; G2 editorial content); Tier 3 — analyst firm content (Forrester, IDC, Ventana Research, G2 Research); Tier 4 — high-quality independent blogs and newsletters with 10,000+ subscribers in your category. The mistake most B2B companies make is pursuing guest posts on Tier 4 publications at scale while ignoring the far harder but far more valuable Tier 1 and 2 placements. A single Harvard Business Review article carries more LLM training data weight than 50 mid-tier blog guest posts, because HBR content is disproportionately represented in the curated, high-quality corpora used by foundation model trainers.

The most scalable path to Tier 1 and 2 mentions is original research. Publish an annual research report on a topic in your category — survey 200–500 professionals, produce original data, and pitch the findings to Tier 1 publications as a story. Original research generates citations from other publications, appears in multiple LLM training sources, and creates a recurring content asset that builds compound authority over time. The research topic should be at the intersection of your buyer's most pressing concern and a question that no one currently has definitive data on. A well-executed annual research report generates 15–40 earned media placements from a single primary publication relationship.

- Tier 1 targets: Forbes, HBR, MIT Sloan, McKinsey Insights, Gartner — prioritize 1-2 placements per year
- Tier 2 targets: category-specific trade publications with editorial standards; target 6-10 per year
- Highest-leverage content type: original annual research report with primary survey data
- Original research target: 200-500 respondents; pitch findings to Tier 1 publications first
- Award and list inclusions: G2 category leader badges, Forrester Wave inclusions — pursue annually
- Conference speaking: industry conference speaker listings appear in LLM training corpora; prioritize 3-4/year

- Track: document every Tier 1/2 placement with publication DA, topic, and month — review breadth quarterly

One original research report published to a Tier 1 outlet generates more LLM training data authority than 12 months of typical guest posting. Invest in primary research — it is the highest-leverage activity in an LLM SEO strategy.

SECTION 6

The Brand Language Consistency System: Same Words, Every Platform

LLMs build entity representations from patterns in text. When the same brand is described as a 'marketing automation platform' on its website, a 'RevOps software tool' on G2, a 'demand generation company' in a press release, and a 'sales enablement solution' in a guest post, the model encounters four different category labels for the same entity. This inconsistency fragments the entity representation and reduces citation confidence — the model is less likely to cite your brand when it is uncertain what category you belong to. The solution is a brand language system: a canonical set of 5–7 key descriptors that are used consistently everywhere. This includes your primary category descriptor (the one-phrase answer to 'what does your company do'), your three core use case descriptions, your one-sentence company description, your differentiating attribute phrase, and your target customer descriptor. These exact phrases should appear on your website's homepage and About page, your LinkedIn company summary, your Crunchbase description, your G2 profile description, your Twitter/X bio, your schema.org description property, and in the briefs you provide to PR firms and guest post editors. Build a 'language lockdown' document that every person who writes about or for your company receives — this is a non-negotiable tool for maintaining entity clarity at scale.

The audit process: extract your brand descriptor from each platform where you have a profile and compare them to each other and to your canonical language document. Inconsistencies appear in three common forms: synonym drift (using different but related words for the same concept), scope creep (describing your product differently for different audiences without core language consistency), and outdated positioning language that was updated on the website but never propagated to third-party profiles. Schedule a cross-platform language audit quarterly and update any profile that deviates from canonical language within 2 weeks of the audit.

- Define 5-7 canonical brand language elements: category, use cases, one-sentence description, differentiator
- Apply canonical language to: website homepage/About, LinkedIn, Crunchbase, G2, schema.org, Twitter/X bio
- Create a 'Language Lockdown' document distributed to anyone writing about or for your company
- Audit all third-party profiles quarterly for language drift; update deviations within 2 weeks
- Provide PR firm and guest post editors with exact brand language in their brief — not a general boilerplate
- Never use different category descriptors for different audiences — consistency beats tailored messaging in entity building

- Test: search your brand in ChatGPT and note how it describes your category — this reveals the model's current entity representation

58%

of B2B brands use three or more different category labels across their public profiles, fragmenting their entity signal

SECTION 7

Measuring LLM Visibility: The 30-Query Tracking Protocol

LLM visibility cannot be measured by traditional rank tracking tools — there are no positions, no SERPs, and no consistent output for the same query across different users and contexts. What can be measured is citation frequency: the percentage of relevant queries where your brand is cited in the LLM's response, measured consistently over time using a controlled query set. The 30-query tracking protocol works as follows. Build a query library of 30 prompts across three categories: (1) Category definition queries — 'What is [your category]?', 'How does [your category] work?', 'What are the benefits of [your category]?'; (2) Recommendation queries — 'What are the best [your category] platforms for [your target customer size]?', 'What should I look for when choosing a [your category] tool?', 'Compare [your category] options for B2B companies'; (3) Problem-solution queries — 'How do I [solve the core problem your product addresses]?', 'What tools help with [specific pain point]?'. Run all 30 queries across ChatGPT (GPT-4o), Perplexity, Google Gemini, and Claude on the first Monday of every month. Log whether your brand is cited in each response (yes/no), whether the citation is positive, neutral, or negative, and which competitor brands are cited in the same response. Calculate your citation rate for each platform: $(\text{number of queries where cited} / 30) \times 100$.

Track your citation rate monthly as a primary KPI. A new or underdeveloped entity will typically start at 0–10% citation rate. After 6 months of systematic entity building and content architecture work, a well-executed program should reach 25–40% citation rate. Above 50% citation rate on recommendation queries for your specific category is excellent and indicates strong entity authority. Also track competitive citation share — what percentage of responses that mention your category cite at least one competitor? This tells you whether your entity authority is keeping pace with your competitive landscape. Run the full 30-query audit, not a shortcut sample — the query diversity matters for catching platform-specific gaps.

- Build 30-query library: 10 category definition, 10 recommendation, 10 problem-solution queries
- Run queries on: ChatGPT (GPT-4o), Perplexity, Google Gemini, Claude — monthly on the same date
- Log for each query: cited (yes/no), sentiment (positive/neutral/negative), competitors cited
- Calculate citation rate per platform: $(\text{citations} / 30) \times 100$
- Track monthly: overall citation rate, citation rate by query type, competitive citation share
- Benchmarks: 0-10% = new entity; 25-40% = developing authority; 50%+ = strong authority
- Review query library quarterly — add new queries as your product and category vocabulary evolve

Run the same queries in incognito/logged-out mode to avoid personalization bias in LLM responses. Log results in a shared spreadsheet immediately — LLM outputs are not persistent and cannot be retroactively retrieved.

SECTION 8

Platform-Specific Tactics: ChatGPT, Perplexity, Gemini, Claude Differences

ChatGPT (without browsing enabled) is primarily training-data-driven in its citations for factual and recommendation queries. The implication is that training data volume and source authority matter most — publication history before the model's training cutoff is the primary lever. Focus on building a strong historical publication footprint in high-authority sources and ensuring your entity is well-defined in sources known to be in major LLM training corpora (Wikipedia, Crunchbase, major trade publications). When ChatGPT browsing is enabled, it behaves more like a search engine — real-time content quality and SEO authority become relevant. Perplexity is almost entirely RAG-based and uses a search index similar to Bing for retrieval. Perplexity optimization is essentially advanced SEO: rank in the top 5 for your target queries in a search engine that Perplexity retrieves from, ensure your content is structured with direct-answer format so it is extractable, and earn high-authority backlinks that signal content trust. Perplexity also allows direct source submission through its add sources feature, which can accelerate your presence in its index.

Google Gemini blends Google's search index with its Knowledge Graph and training data. The Knowledge Graph influence means entity setup — particularly Google Business Profile, schema.org markup, and Wikipedia/Wikidata — carries extra weight compared to other platforms. Strong Google Search rankings for target queries also directly influence Gemini's retrieval. Claude (Anthropic) is primarily training-data-based and shows a notable preference for sources that demonstrate structured expertise — long-form authoritative content, academic or research-style writing, and factually well-sourced pieces. Claude's Constitutional AI training gives higher weight to accurate, well-evidenced content over promotional language. Write for Claude by prioritizing factual depth, specific data points, and clear logical structure over persuasive or marketing-oriented framing.

- ChatGPT (no browsing): prioritize historical publication volume in Tier 1-2 sources + entity completeness
- ChatGPT (with browsing): optimize for search rankings on target queries; same as Perplexity tactics
- Perplexity: advanced SEO for Bing index; top-5 rankings; direct-answer content structure; consider source submission
- Gemini: entity setup is especially high-leverage; Google Knowledge Graph + Business Profile + schema.org
- Claude: factual depth and well-evidenced claims outperform marketing framing; cite your own data and research

- Cross-platform priority: direct-answer content format and entity completeness benefit all four platforms
- Platform gap analysis: if you are cited on 3 platforms but not 1, use the 30-query log to diagnose which signal is missing

2.8x

higher Perplexity citation rate for brands ranking in the top 3 Google results for their category query vs. brands ranking 4-10

SECTION 9

6-Month LLM SEO Roadmap

Month 1 — Entity foundation: Implement schema.org Organization markup with sameAs array. Complete Crunchbase, G2/Capterra, LinkedIn, and Wikidata profiles using canonical brand language. Run initial 30-query audit to establish baseline citation rates across all four platforms. Document current citation rate by platform and query type. Month 2 — Content architecture: Audit existing content for direct-answer format compliance. Identify and rewrite the 5 highest-traffic category pages to direct-answer structure. Publish 2–4 new direct-answer content pieces targeting core category definition and recommendation queries. Month 3 — Publication outreach: Launch an original research initiative (survey design and distribution). Simultaneously pitch 3–5 byline placement opportunities to Tier 2 industry publications. Begin building a media list of 20+ journalists and editors who cover your category. Month 4 — Research publication and amplification: Publish the original research report. Pitch the data to Tier 1 publications and issue a press release with key findings. This is your highest-impact authority-building month.

Month 5 — Consistency audit and expansion: Conduct cross-platform language audit; update all profiles with canonical language. Submit to any analyst firm lists or award programs relevant to your category. Add at least 10 new external entity references (press mentions, directory listings, event speaker profiles). Month 6 — Measurement and optimization: Run the full 30-query audit again. Compare Month 6 citation rates against Month 1 baseline. For any platform where citation rate has not improved, identify the specific signal gap (entity, content format, or authority breadth) and adjust the plan. Commit to a quarterly publication of the original research to maintain authority compound growth.

- Month 1: Schema.org + Wikidata + G2 + Crunchbase + LinkedIn setup; baseline 30-query audit
- Month 2: Direct-answer content audit; rewrite 5 core category pages; publish 2-4 new pieces
- Month 3: Original research survey launch; Tier 2 byline pitches; media list build
- Month 4: Research report publication; Tier 1 pitch; press release distribution
- Month 5: Cross-platform language audit; analyst/award submissions; 10+ new external mentions
- Month 6: Full 30-query audit; compare to baseline; identify platform-specific gaps; quarterly plan

Month 4 is the inflection point. Original research publication combined with Tier 1 pitching creates the compounding authority effect that separates brands that plateau at 15% citation rates from those that reach 40%+.

Implementation Checklist

Phase 1 — Entity Foundation

- Implement schema.org Organization markup with all required properties including sameAs array
- Complete Crunchbase profile with canonical brand language in every field
- Complete G2 or Capterra profile; verify category selection matches schema.org description
- Create or update Wikidata entity with all structured properties and external ID links
- Verify Google Knowledge Panel appears for brand search; request corrections if inaccurate
- Conduct cross-platform language audit; standardize all profiles to canonical brand descriptors
- Run initial 30-query audit across ChatGPT, Perplexity, Gemini, Claude; document baseline

Phase 2 — Content and Authority Build

- Audit top 10 content pages for direct-answer format compliance; rewrite non-compliant pages
- Create direct-answer category landing pages for top 6-8 buyer research questions
- Design and launch original research survey (200+ respondents, category-relevant topic)
- Pitch 3-5 Tier 2 industry publication byline placements using original research data
- Pitch 1-2 Tier 1 publication placements (Forbes, HBR, McKinsey) with research findings
- Submit to 2-3 analyst firm reports or category list inclusions (G2 Leader, Forrester Wave)
- Build speaker listing presence at 2-4 industry conferences in your category

Phase 3 — Measure and Optimize

- Run 30-query audit monthly; track citation rate by platform and query type

- Quarterly: cross-platform language audit; update any drifted profile descriptions

- Quarterly: publish updated research report or new original research to maintain authority velocity

- Identify platform-specific citation gaps in monthly audit; match gap to signal type (entity/content/authority)

- Expand 30-query library annually as product and category vocabulary evolves

NetWebMedia

We Build LLM Citation Strategies for B2B Brands

NetWebMedia develops and executes LLM SEO programs for B2B companies — from entity authority setup and content architecture restructuring to original research strategy and high-authority publication placement. We measure citation rates monthly and optimize based on platform-specific performance data.

AI Marketing Automation

AEO & AI-First SEO

Autonomous AI Agents

Paid Media + AI Creative

CRM + AI Workflows

netwebmedia.com/contact