

BRAINPAN.AI · STRATEGIC BRIEF

The Business Case for AI Visibility

A boardroom-ready argument for CMOs, CDOs, and marketing leaders, grounded first in Brainpan's own audited AI VI™ benchmark data and secondarily in fact-checked, source-attributed industry evidence.

134

Brands audited in our own benchmark

1,742

Organic citations analyzed

5

AI assistants measured

EXECUTIVE SUMMARY

Three things every CMO needs to know

01**AI search now shapes B2B purchase decisions.**

Buyers use ChatGPT, Gemini, Perplexity, Copilot, and Claude to identify vendors and shortlist before speaking to sales. [Forrester reports](#) that 94% of business buyers now use AI at some point during their buying process. Brands named in those responses have a structural advantage; brands absent face a gap traditional SEO metrics cannot see.

02**Strong SEO does not guarantee AI citation.**

In our own Q2 2026 audit of 134 insurance brands, two of five AI assistants tested (Perplexity and Copilot) accounted for 100% of the 1,742 organic citations recorded. ChatGPT, Gemini, and Claude produced mentions but effectively zero organic citations in that category. Existing organic investment does not automatically transfer to AI citation.

03**The baseline is measurable from day one. ROI compounds after.**

Share-of-model and citation baselines are established within 5 to 10 days using the same methodology behind our AIVI™ benchmark. Commercial ROI (AI-referred traffic, assisted conversions, influenced pipeline) becomes reportable as it accumulates, typically over a 60 to 180 day window that depends on your sales cycle and deal volume.

BRAINPAN.AI FIRST-PARTY DATA

A note on the evidence in this brief

Most AI-visibility business-case material leans on vendor case studies that are difficult to verify. We checked every third-party statistic against its original source, correcting several and removing any we could not source, and we lead with our own audited AIVI™ benchmark data wherever possible. Corrections are noted inline and logged on the web version of this brief.

Decision this brief supports: whether to fund a baseline audit of how your brand is mentioned, recommended, and cited across the AI assistants your buyers use, before committing to a larger program.

OUR OWN EVIDENCE

The AIVI™ Q2 2026 Insurance Benchmark

BRAINPAN.AI FIRST-PARTY DATA

What we found when we audited a real category ourselves

We measured 134 U.S. insurance brands across five AI assistants (ChatGPT, Gemini, Claude, Perplexity, and Copilot) using 300 stratified prompts, producing 1,500 model responses and 4,328 organic brand mentions. Every score is organic-only; paid placements were excluded and disclosed separately.

134

Insurance brands audited

300

Stratified prompts, 5 AI engines

4,328

Organic brand mentions recorded

1,742

Organic citations, by source domain

100%

Perplexity citation rate (925 mentions)

95.7%

Copilot citation rate (854 mentions, 65.6% Top-3)

The finding that should reset expectations: all 1,742 organic citations in this study came from Perplexity or Copilot. ChatGPT, Gemini, and Claude mentioned brands but rarely or never attached a clickable, source-attributed citation in this category during the audit window.

Of the 1,742 citations, 1,348 resolved to a brand within the scored cohort and the remaining 394 resolved to sources outside it.

This is why we score AIVI™ on a weighted five-component index (Share of Model 35%, Recommendation strength 25%, Position weighting 20%, Citation Influence 15%, Top-3 rate 5%) rather than a single citation-count metric.

Full methodology and per-brand results: brainpan.ai/research/aivi-insurance-benchmark-q2-2026/ and brainpan.ai/research/aivi-methodology/. This is a single-category study; other categories may behave differently.

SUPPORTING INDUSTRY SIGNALS

What vendor-reported case studies add to the picture

These are third-party, vendor-published figures, not Brainpan-audited data. We checked each against its original source, linked beside each figure (see Industry evidence, next page). Figures whose source we could not locate are not included. Treat them as directional evidence, not benchmarks to plan against.

3.2×

Increase in AI search conversions within 6 months **VENDOR-REPORTED**

SteelSeries, as reported in a [Goodie case study](#) (six-month window). NoGood's own [write-up](#) of the same work reports different figures (27x year over year in its headline, 7.1x in its body), so this is unreconciled.

458%

Growth in AI Overview visibility for a home services franchise group **VENDOR-REPORTED**

[Ignite Visibility case study](#). The source gives no period, baseline, or measurement method.

All results on this page are drawn from published third-party case studies and presented as directional evidence. They are not guarantees, and not a substitute for your own audited baseline.

INDUSTRY EVIDENCE

Case studies, checked against their sources

B2B SaaS**Presentation software platform (Mentimeter)**

Siege Media reported 124K ChatGPT-referred sessions over a six-month window and 3,400 reported conversions in a separate one-month window. The two figures cover different periods and do not combine into a conversion rate, and the source does not state that all 3,400 were AI-attributed. An earlier "2.7%" figure was removed.

[Siege Media case study](#) · Vendor-reported

Cybersecurity enterprise**Global network security vendor (Fortinet)**

LeadWalnut reports a starting AI citation share of 0.6% (ranked #31) in cloud security as of October 3, 2025, for a company with 700K+ customers. A large customer base does not automatically translate into AI citation. The figure comes from LeadWalnut's own case study of its work for Fortinet, which gives conflicting timelines for later results, so we cite no improvement figures.

[LeadWalnut case study](#) · Vendor-reported

Ecommerce**Premium cookware brand (DI ORO)**

AEO Engine reports +559% organic sales growth and +306% organic traffic growth over four months for this client. The source lists the figures without describing the tactics and does not isolate how much is attributable to AI-referred traffic.

[AEO Engine client results](#) · Vendor-reported

MEASUREMENT

The AI visibility ROI framework

We measure AI visibility using the same five-component AIVI™ taxonomy behind our benchmark, plus downstream performance metrics. Denominators matter: "share" numbers are only comparable when the query set and response count are stated.

Share of Model (35% of AIVI™)

Your brand's share of all eligible organic brand mentions across the competitive set. Mention rate, the percentage of eligible responses that name your brand at least once, is a related but separate measure. Neither equals being cited or recommended.

Citation share and Influence (15%)

Of the responses that mention you, what share include a clickable, source-attributed citation, and how influential is it? This varies widely by assistant.

Recommendation and Top-3 rate (25% + 5%)

How often you are positioned as a recommended option, and how often you land in the top 3 of a ranked response, each stated against the response count measured.

AI-referred performance

Sessions and conversions from AI referrals. GA4 and Adobe Analytics each provide an AI-assistant referral classification under different names, so confirm how yours is configured; traffic with no referrer stays unattributed. Report assisted and last-click separately.

Volatility check: in [Adobe's U.S. retail data](#), AI-referred visits converted to purchases 42% better than non-AI visits in March 2026, versus 38% worse in March 2025. This is a retail measure, not evidence of B2B lead or pipeline lift.

FOR THE CFO: HOW TO FRAME THE ROI MATH
$$\text{ROI} = (\text{Incremental gross profit attributable to the program} - \text{Program cost}) / \text{Program cost}$$
$$\text{Expected pipeline value (scenario)} = \text{Influenced pipeline} \times \text{win rate} \times \text{gross margin}$$

State your attribution confidence explicitly; influence is not incrementality. The baseline is measurable from day one. Commercial ROI accumulates as AI-referred traffic, assisted conversions, and influenced pipeline build, typically reportable within a 60 to 180 day window that depends on sales-cycle length and deal volume.

MEASUREMENT

Illustrative program timeline

Days 1 to 10**AI Visibility Audit delivered**

Written report. Citation baseline across up to eight AI platforms, scored on AIVI™. Competitive gap map. Corrected structured-data templates. Prioritized 90-day roadmap.

Month 1**Foundation layer**

Schema corrections deployed for entity clarity and machine readability. AEO content restructuring begins. Low-cost hygiene steps such as publishing llms.txt may be added; evidence that it reliably increases citation likelihood is weak, so we treat it as optional, not a headline lever.

Months 2 to 3**Entity clarity and content build**

FAQ and HowTo structured data added where relevant for entity clarity. Google fully removed HowTo rich results (September 2023) and FAQ rich results (May 2026), so this is a clarity step for AI systems, not a guaranteed citation tactic. A/B testing of entity citations begins on third-party sources (G2, Crunchbase, Wikipedia, blog hubs).

Months 4 to 6**Share-of-model growth**

Published vendor case studies report citation and visibility growth ranging widely by program. AI-referred sessions begin appearing in analytics, and AI visibility often shows up as direct-traffic and branded-search lift even without a clickable referral ("dark" AI-influenced traffic).

Months 7 to 12**Compounding authority**

Established source authority and broad third-party corroboration may make competitors harder to displace over time. That is a plausible mechanism, not a proven guarantee. AI-referred and influenced-pipeline data becomes reportable for CFO review.

Timeline is illustrative. Actual results depend on starting position, content volume, category competitiveness, and execution velocity.

INTERNAL BUY-IN

The argument by stakeholder

For the CEO / Board

[Forrester reports](#) that 94% of business buyers now use AI at some point during their buying process (up from roughly 89% a year earlier). Brands absent from AI-generated responses risk losing consideration-set inclusion before the first sales conversation.

For the CFO

AI visibility produces measurable outputs: share-of-model improvement from an audited baseline, AI-referred traffic in existing analytics, and pipeline influence reportable against program cost using $ROI = (\text{incremental gross profit} - \text{program cost}) / \text{program cost}$.

For the CRO / VP Sales

AI-referred lead quality is promising but volatile: in [Adobe's U.S. retail dataset](#), AI-referred visits converted to purchases 42% better than non-AI visits in March 2026, after converting 38% worse in March 2025. That is a retail measure, so treat it as a signal to test in B2B, not a lead-quality benchmark.

For the CMO

AI visibility integrates with and amplifies existing SEO and content investment: structured data, content architecture, entity clarity, and third-party citation building on sites like G2 and Wikipedia.

For the CDO

AI visibility depends on systems and data a CDO owns: who controls structured data and CMS changes, how AI-referred traffic and self-reported discovery join analytics and CRM reporting, and what measurement governance and consent rules apply. A baseline audit gives the data team a defined scope to decide what to instrument, who owns schema changes, and how results connect to pipeline reporting.

The cost of inaction likely compounds. Established source authority and broad third-party corroboration may make competitors harder to displace once they have built citation presence. That is a plausible mechanism, not a proven guarantee. Brands that establish an audited baseline now are, at minimum, not flying blind on a channel already influencing 94% of B2B buying processes.



Ready to see your citation footprint?

The AI Visibility Audit is a written diagnostic delivered within 5 to 10 business days. It establishes your real citation baseline across up to eight AI platforms, including ChatGPT, Gemini, Perplexity, Copilot, Claude, and Google AI Overviews, scored on the same AIVI™ methodology behind our published benchmark, with competitive gap analysis, schema recommendations, and a prioritized 90-day roadmap.

Written audit report (25 to 47 pages)

Corrected structured-data templates

90-day action roadmap

No introductory sales call required

Request AI Visibility Audit

brainpan.ai/contact/ · kwalsh@brainpan.ai · [+1 703-951-7195](tel:+17039517195)

APPENDIX

Sources & references

First-party data

1. Brainpan.AI, AIVI™ Q2 2026 Insurance Benchmark. brainpan.ai/research/aivi-insurance-benchmark-q2-2026/
2. Brainpan.AI, AIVI™ scoring methodology. brainpan.ai/research/aivi-methodology/

Third-party statistics (checked against their sources)

1. Forrester, [B2B Buyers Make Zero-Click Buying Number One](#) (94% of business buyers use AI during the buying process, up from 89%).
2. Adobe, [U.S. retail AI traffic report, April 2026](#) (AI-referred visits converting to purchases 42% better in March 2026, 38% worse in March 2025; U.S. retail sites).
3. Goodie, [SteelSeries case study](#) (3.2× AI search conversion increase over six months; vendor-reported). NoGood, [SteelSeries case study](#) (different figures: 27x year over year, 7.1x in body text).
4. Ignite Visibility, [home services franchise AEO case study](#) (458% increase in AI Overview visibility; period not stated).
5. [Siege Media](#), Mentimeter case study (124K ChatGPT sessions over six months; 3.4K reported conversions in one month; non-overlapping periods).
6. LeadWalnut, [Fortinet case study](#) (starting citation share 0.6%, ranked #31, cloud security, October 3, 2025; vendor-reported).
7. AEO Engine, [client results](#): DI ORO (+559% organic sales, +306% organic traffic, four months; vendor-reported).

Search feature status (verified)

1. Google removed HowTo rich results from search (mobile and desktop), September 13, 2023.
2. Google removed FAQ rich results from search (universal), May 7, 2026.

Report a correction: kwalsh@brainpan.ai