



GENERATIVE SEARCH

CONTENT STRATEGY

B2B EDITORIAL

Winning Search in the Age of *Generative Discovery*

A strategy framework for Winsome Marketing clients – what's changed, why, and how to win the next era of organic visibility.

Vol.
01

CREATING PHENOMENA.



INSIDE THIS ISSUE

A field guide to organic visibility, post-AI Mode.

01	What Google Announced — and what it actually means	P. 03
02	The zero-click reality and what it means for lead quality	P. 04
03	GEO research findings you should know	P. 05
04	The CORE finding: content positioning isn't random	P. 06
05	Earned media: why brand-owned content alone isn't enough	P. 07
06	How different AI engines use different sources	P. 08
07	Winsome's editorial framework	P. 09
08	Case study · A top-100 accounting firm	P. 10
09	The consumer shift & what to do next	P. 11

A NOTE FROM WINSOME

//

Search has changed more in six months than it did in the previous six years.

This document explains what's actually happening inside generative engines, what the research now confirms, and how to upgrade an editorial program so it's cited — not just indexed — in 2026 and beyond.

RESEARCH-LED

FOR WINSOME CLIENTS

PUBLICATION

Volume 01

01 · THE SHIFT

What Google announced – and what it actually means.

Google's latest I/O announcement confirmed several things that matter for how organic search works. The most important is the simplest: AI Mode is no longer a test.

It's the default for a billion people, powered by Gemini 3.5 Flash, with a redesigned search box that no longer assumes keywords – it assumes natural-language questions, often long, often complex, often submitted alongside images, files, or browser tabs. Users ask follow-up questions directly from AI Overviews and flow into conversational threads. Context persists across that thread.

This is a fundamentally different discovery environment than a ten-blue-links SERP. When a user asks a complex question conversationally, the engine doesn't return a list of pages and let them decide. It synthesizes. It picks sources, extracts information, and assembles an answer – with links surfaced as supporting evidence, not as the primary result.

Additionally, Google announced **information agents** – AI that monitors topics 24/7 on a user's behalf, synthesizes updates, and delivers them proactively. That's not a search the user initiated. That's your content being retrieved and evaluated while the user is doing something else entirely.

AI MODE USERS

1B_{mo.}

Google crossed one billion monthly users on AI Mode in 2026.

QUERY VELOCITY

2_x/Q

AI Mode queries have doubled every quarter since launch.

THE IMPLICATION

Search volume keeps rising – but the path from query to website is getting more attenuated. The content that does earn a click is worth more, and being the one cited matters more.

ANATOMY OF AN AI MODE ANSWER

QUERY
Long & conversational

SYNTHESIS
Gemini 3.5
picks · extracts
· assembles

CITED #1

CITED #2

CITED #3

02 · THE ZERO-CLICK REALITY

The clicks that *do* get through are worth more.

The concern Winsome clients most often raise about AI search is the zero-click problem. If AI Overviews answer the question at the SERP level, the visitor never reaches the site. The concern is legitimate — for sites built heavily around informational traffic, the click decline is real and documented.

But the data on what happens to the clicks that *do* get through tells a more interesting story.

LLM VISITORS

4.4x

downstream value vs. average organic-search visitor.

AI SEARCH VISITORS

23x

higher conversion rate than standard organic — at a small fraction of total sessions.

A traditional organic visitor is still in *exploration* mode — running queries, reading sites, forming a picture. A visitor arriving from an AI tool has already had the exploratory conversation inside the AI platform. They've asked follow-ups, compared options, and gotten summaries. By the time they click, they're much closer to a decision.

REFRAME

The AI handled the research. The site handles the conversion. Both jobs matter — and conflating them is a common mistake.

VISITOR JOURNEY, TWO PATHS

Where the research happens determines lead quality.

TRADITIONAL ORGANIC



Research happens *on your site* across many pages. Intent matures slowly. Conversion is far downstream.

AI-MEDIATED



Research happens *inside the AI tool*. The click lands a decision-ready buyer on a single page.

For clients whose goal is qualified lead generation rather than traffic volume, this is a favourable shift. A smaller number of highly qualified visitors who are close to a purchase decision is more valuable than large volumes of informational traffic with low conversion potential.

03 · WHAT THE RESEARCH SAYS

GEO findings you should know.

Aggarwal et al., KDD 2024 – 10,000 queries across generative engines. Some interventions reliably improved citation rates. Some actively hurt them.

VISIBILITY LIFT BY INTERVENTION

Citation-rate change against control. Negative is worse than doing nothing.

Verifiable statistics		+40%
Direct quotations		+22%
Citing external sources		+17%
Fluency + specificity best-performing pair		+45%
Keyword stuffing classic "SEO tactic"		-10%

Source: Aggarwal et al., "Generative Engine Optimization," KDD 2024. Lift figures rounded for clarity; "fluency + specificity" reported as >5.5% over best single intervention.

There's a lot of noise in the market right now about generative engine optimization — agencies rebranding classic SEO as something new, case studies that don't survive scrutiny, and frameworks that amount to keyword stuffing dressed up in AI language. The research tells a more specific and more useful story.

The signal generative engines use to decide what to cite has almost nothing to do with keyword density and almost everything to do with whether the content provides clear, verifiable, well-attributed answers.

THE PRACTICAL TRANSLATION

- At least one specific, verifiable statistic anchored to a citable source per piece.
- At least one direct claim extractable as an answer to the keyword's implied question.
- A structure that surfaces the relevant claim quickly — not buried mid-paragraph.

THE UNCOMFORTABLE TRUTH

The single most widely-practiced SEO tactic — keyword density — performs worse than doing nothing in generative engines.

These are not stylistic choices. They are the structural properties that determine whether your content gets cited or passed over.

04 · THE CORE FINDING

Generative ranking *isn't* random — and content shape is the lever.

Jin et al. (University of Illinois, Feb 2026) studied how GPT-4o, Gemini 2.5, Claude, and Grok actually rank content when synthesizing answers. Two content shapes dramatically outperformed everything else.

SHAPE A

Reasoning-style content

TOP-1 PROMOTION RATE

80%+

Walks through a logical argument, compares alternatives, and arrives at a well-supported conclusion. Elevated to Top-1 position at rates above 80% across all four models tested.

SHAPE B

Review-style content

DETECTABILITY ABOVE BASELINE

~ O Δ

First-person, narrative, reflects lived experience and genuine comparison. Detection rates by human evaluators only slightly above baseline. Especially effective on Gemini and Grok.

//

The structural qualities that help AI systems trust and cite your content are the same qualities that make content genuinely useful to human readers.


Content written to *explain*, to *compare*, to *narrate experience*, and to *walk through reasoning* performs better in generative retrieval than content written purely to rank for keywords.

This is not a loophole. It's confirmation that genuine topical authority — expressed through well-structured, specific, experience-grounded content — is the durable position in both traditional and generative search.

05 · EARNED MEDIA

AI engines weight your *third-party footprint* above your own site.

CITATION MIX · CLAUDE · CONSUMER ELECTRONICS QUERIES



Claude's consumer-electronics results were composed of approximately **93.7% earned media**. ChatGPT runs close behind. This isn't a minor preference — it's the architecture of how these systems decide what to trust.

One of the most consistent findings across multiple independent studies is that AI search engines heavily favor what researchers call **earned media** — third-party coverage, editorial mentions, expert reviews, institutional citations — over brand-owned content.

A 2025 study analyzing how AI systems assemble search responses across Claude, ChatGPT, Perplexity, and Gemini found that earned content dominated across all engines at all query intent levels. For informational queries, earned sources accounted for the vast majority of citations. Even at the transactional end — when someone is essentially asking "which product should I buy" — brand-owned content competed with, but rarely dominated, third-party sources.

For a professional services firm, this creates a specific implication: your website content, no matter how well-structured and keyword-optimized, operates in a system that weights it *below* third-party coverage of your firm and your expertise.

The response isn't to abandon your own site — it's foundational, and the research is clear that you must be indexed and trusted in traditional search before generative engines can cite you at all. But the editorial program has to be paired with a systematic effort to build the third-party coverage that AI systems treat as verification.

STRATEGIC IMPLICATION

The content program and the PR program are the same program — not separate workstreams with separate goals.


06 · ENGINE MIX

Not all AI engines work the same way.

Treating them as equivalent is a strategic mistake. Visibility tactics that work on Gemini quietly under-perform on ChatGPT and Claude.

APPROXIMATE SOURCE WEIGHT BY ENGINE

EARNED BRAND SOCIAL



Indicative composition based on 2025 cross-engine analyses. Exact mix shifts by vertical and query intent.

CHATGPT

Earned-heavy

High cross-language stability. Draws on the same authoritative third-party sources regardless of phrasing.

CLAUDE

Most earned-heavy

~94% earned on consumer electronics queries. Trust is built almost entirely off-site.

PERPLEXITY

Balanced mix

Sensitive to YouTube, Reddit, community discussions. Trade-press coverage carries it forward.

GEMINI

Brand-friendliest

Incorporates more brand-owned properties. Well-structured site content carries more weight here.

A strategy that focuses only on Google rankings and your own site will perform reasonably on Gemini, inconsistently on Perplexity, and poorly on ChatGPT and Claude for anything other than direct brand queries — because you haven't built the third-party authority those systems require.

07 · THE FRAMEWORK

Winsome's editorial framework — five upgrades, no pivot.

*Topical authority through volume, structure, and intent alignment still works.
The CORE research adds quality constraints — not a new playbook.*

01

Structure content to be *extracted*, not just read.

Every article should have a clearly identifiable answer to the question the keyword implies. Direct answers near the top. H2s that frame a position, not just a topic. FAQs that mirror conversational follow-ups. Structured comparison language — "X works better than Y when..." — rather than balanced non-statements.

02

Shift from informational coverage to expert narrative.

Review-style content sounds like someone who has actually done the thing being described. Activate subject-matter experts in the program. For B2B fields where the advice is high-stakes — tax, accounting, legal, finance — this is the competitive moat. Generative engines are learning to tell the difference.

03

Match content volume to topical depth — not quantity targets.

More articles with genuine depth and specific angles outperform the same ground at lower specificity. The editorial calendar question isn't "how many posts" — it's "which specific questions within our topic clusters are still unanswered well enough to get cited."

04

Use automated production as a floor, not a ceiling.

AI handles research, first draft, and structural scaffolding at pace. A human editor with subject-matter access adds the reasoning layer — the why, the comparison, the specific insight nobody else could have produced. That's the production infrastructure that makes topical authority scalable.

05

Track generative visibility — not just rankings & sessions.

AI Overview appearances in GSC. Direct traffic correlated with publishing windows. Brand-mention volume across AI tools. Emerging AI referral sources. The goal is a complete picture of generative discoverability, not just the portion of it that looks like a click.

08 · IN THE FIELD

Case study — a top-100 accounting firm.

Repositioning as an AI-native practice. CFOs and controllers at \$10M–\$100M companies. Six months in.

ORGANIC CLICKS

+ 87.9%

5,990 → 11,260 monthly Google organic clicks.

SEARCH IMPRESSIONS

+ 169%

839K → 2.26M impressions.

RANKING KEYWORDS

+ 58%

8,225 → 13,000+ total.

AI OVERVIEW VISIBILITY

577 kw

At "strong authority" threshold for professional services.

GENERATIVE ENGINE TRAFFIC

54.7%

Engagement rate — above direct traffic benchmark.

MONTHLY AI AUDIENCE

12.8M

Across all generative engines, tracked.

CHATGPT-REFERRED DEAL · CLOSED

\$112,400

Construction vertical.

CHATGPT-REFERRED DEAL · CLOSED

\$84,471

Construction vertical.

AI SENTIMENT RANK


1st or 2nd

In five of seven competitive categories tracked across generative engines.

THE MECHANICS

A consistent high-volume editorial cadence anchored to vertical-specific expertise; a newsjacking program that put the firm's voice into active news cycles; systematic PR activation expanding the backlink profile from Tier 1–3 earned media; and a structured content approach where every article contains the kind of expert-grounded, comparison-style reasoning that both traditional algorithms and generative systems are trained to prefer.

THE SPLIT



09 · THE CONSUMER SHIFT

People are no longer searching to find a list.

They're searching to get an answer. They ask longer, more specific questions. They follow up. They let AI synthesize information they used to compare themselves. The search box has become a conversation partner, not a directory.

That shift was already underway before Google announced it — AI Mode's billion users didn't appear overnight. What the announcement confirmed is that this behaviour is now the default, and the SERP is being redesigned around it.

//

The content that earns the answer earns the relationship. Generic coverage earns nothing — because generic coverage gets synthesized away.

WHAT TO DO NEXT

Two questions worth answering honestly about your existing content.

01

Is it structured so that a generative engine could extract a clear, specific answer from it?

02

Does it contain the kind of expert-grounded specificity that would make an AI system choose it over competitors' coverage?

If the answer to either is *no*, that's where the program evolves. Winsome can run that audit, identify the highest-leverage gaps, and build the editorial and production workflow to close them.



THE NEXT MOVE

The methodology exists.
The tools exist. The only
question is how much
generative search share you
want to own before your
competitors figure out the
same thing.

START A CONVERSATION

REQUEST AN AUDIT

WINSOME MARKETING

Creating phenomena.

Strategy-led, content-first growth for B2B leaders who want to make
more than an impression — they want results.

VOLUME 01

Winning Search in the Age of Generative
Discovery

WINSOMEMARKETING.COM