



Whitepaper

Commanding Generative Search



→ How Brands Take Back Control in the AI-Defined Internet

(An answer-first playbook for visibility, accuracy, and measurable revenue impact)

Version: August 2025

Audience: CMOs, Growth & Content Leaders, RevOps, Product Marketing, SEO Leads, Data & Analytics, Legal/Comms

Author: Jetstorm (methods abstracted; implementation details withheld by design)

Executive Summary

Generative engines are no longer a novelty layer on top of search—they are becoming the place where decisions are made. Buyers ask one question and get one synthesized answer that compresses research across the open web. In this environment, the fundamental growth question is not “Are we ranking?” but “Are we named and cited” inside the answer that buyers actually read?

This paper codifies a measurement and operating model for the answer economy:

- **Measure what matters:** Answer Share of Voice (A-SOV), Citation Quality Index (CQI), Competitive Salience, and Narrative Alignment by query theme and model (ChatGPT/Google AI/Perplexity).
- **Optimize the surface that the models read:** Proof density, verifiable facts, factual consistency across the web, and authority coverage are tunable inputs that models use to select and frame brands in synthesized answers.
- **Attribute without the click:** Combine tagged answer-to-site journeys, message-match fingerprinting, geo/time-boxed market experiments, and multi-touch modeling to connect “answer presence” to consideration and pipeline.
- **Operate with governance:** Treat answer optimization as a program (not a sprint) with quarterly scorecards, model-specific QA, and legal/comms guardrails—especially across languages and regions.

Mandate: If you're not in the answer, you're not in the consideration set. Jetstorm equips brands to see, shape, and scale presence inside answers—moving from “sometimes cited” to “consistently preferred.”

1: The Generative Answer Economy (Why this changed and why it matters)

What changed

- **Synthesis over lists.** Major engines are increasing the share of queries that show AI answers (e.g., Google's AI Overviews), though frequency varies by intent and market. Multiple large-scale studies observed prevalence in the mid-teens to high-teens of tracked U.S. queries after an initial spike in early tests.
- **Clicks compress.** Independent analyses indicate ~60% of Google searches end without a click, with ~36% of clicks going to the open web—behavior consistent with answer-heavy result pages.
- **Answer engines are scaling.** Perplexity, a citation-forward answer engine, reports ~22M monthly active users and a run-rate that more than doubled across 2024–2025.
- **Budget follows behavior.** Analyst outlooks put 2025 generative AI spend in the hundreds of billions and the broader AI market on track to exceed \$600B within a few years—fueling more answer-first UX.
- **Authority signals matter.** When AI Overviews cite sources, they disproportionately favor high-authority domains; more than half of cited sources in studies were already top-10 organic pages.

Implication: “Rankings” are no longer a reliable proxy for inclusion in the answer. The optimization surface is broader (facts + authority + consensus + structure), the measurement model is different (A-SOV, CQI), and the analytics stack must account for zero-click behaviors.

2 : Definitions & Core Metrics (What to track and how to compute it)

2.1 Answer Share of Voice (A-SOV)

Definition: The percentage of model-answers (for a defined set of buyer prompts) in which your brand is mentioned or cited—by model, market, and language.

- **Scope:** Query Themes (e.g., “best for [use case],” “alternatives to [X],” “compare [X vs Y],” “implementation of [category]”), segmented by funnel stage and intent.
- **Unit of analysis:** (Prompt × Model) pairs.
- **Formula:**
$$\text{A-SOV} = \frac{\text{\# of (Prompt × Model) pairs where brand is present}}{\text{Total \# of (Prompt × Model) pairs}} \times 100$$

Worked Example (toy): You test 12 prompts across 3 models (36 pairs). Your brand appears in 14 answers. A-SOV = $14 / 36 = 38.9\%$.

2.2 Competitive Salience

Definition: Frequency and context in which named competitors appear alongside your brand in the same answer (e.g., neutral mentions vs. recommended over you).

Use: Detect over-recommendation, negative framing, or missing differentiators.

2.3 Narrative Alignment

Definition: Degree to which the synthesized description reflects your current positioning (ICP, value drivers, pricing tier, integrations, deployment model).

Scoring rubric (illustrative): 0 (inaccurate) → 5 (precisely aligned with current messaging).

Note: Jetstorm’s internal classifiers operationalize this rubric without exposing proprietary prompts or labeling methods.

3 : Evidence Snapshot (for buyers and CFOs)

Short, citable facts you can share upstream:

- **AI answer prevalence:** Multiple independent trackers observed AI Overviews appearing for ~15–19% of studied U.S. queries after the early SGE phase, with prevalence varying by query length and industry.
- **Zero-click is the norm:** ~60% of Google searches end without a click; ~36% of clicks go to the open web.
- **Answer engines are mainstreaming:** Perplexity reports ~22M MAU and a rapidly growing query share.
- **Spending is durable:** Analysts forecast ~\$644B in gen-AI spending in 2025; broader AI to ~\$631B by 2028.
- **Authority bias in citations:** Studies show AI Overviews heavily cite high-authority domains; >50% of cited sources already rank top-10.

(**Note:** Perplexity, Copilot/Bing, and Claude all expose citation features or guidance, reinforcing the importance of source authority.*)

4 : Measurement System (How to see what models actually say)

4.1 Inventory the “real buyer” prompt space

- **Theme map:** Buying jobs → prompts buyers actually ask (problem, category, comparative, implementation, ROI, integration, localization).
- **Locales:** Don't stop at English. Prioritize languages/regions by pipeline potential and partner influence.

4.2 Collect the answers (with governance)

- **Sources:** Log model responses via permitted interfaces; capture Google AI Overview snapshots; record Perplexity/Copilot/Claude answers with their cited sources.
- **Compliance:** Respect rate limits, terms, regional data rules, and robots directives. (Jetstorm's program enforces collection guardrails.)

4.3 Score the answers

- **Presence:** Brand mentioned? Cited? Promoted as “best for” vs. “consider”?
- **CQI:** Weight co-occurring citations by authority tier.
- **Narrative:** Apply the alignment rubric to key value claims and ICP descriptors.
- **Sentiment context:** Neutral vs. positive recommendation; flag negative/incorrect claims for remediation.

4.4 Instrument lift tests

- **Geo/time-boxed experiments:** Select matched markets; introduce optimization interventions; read A-SOV/CQI weekly.
- **Outcome mapping:** Correlate A-SOV/CQI shifts with branded queries, direct/organic landings, and opportunity creation (CRM).
- **Design tip:** Use diff-in-diff to isolate impact (e.g., net +14 pp A-SOV lift in test vs. control when control barely moves).

Deliverable: A quarterly “Answer Scorecard” by theme, model, and market with red/yellow/green thresholds tied to pipeline targets.

5: What Engines Reward (Principles—not tactics)

Goal: Provide practical guardrails without revealing Jetstorm's proprietary methods.

Content & evidence principles

1. **Proof density:** Concentrate verifiable facts (specs, versions, prices, compatibility, SLAs, references) in scannable blocks (tables, bullets, FAQs).
2. **Consistency graph:** Keep facts synced across docs, marketplace listings, partner pages, and translated/localized content.
3. **Authority coverage:** Earn and maintain citations from Tier-1/Tier-2 sources (analysts, standards bodies, top-tier media, high-signal communities); avoid over-reliance on low-trust sources.
4. **Structure for retrieval:** Passage-level retrieval favors clean headings, stable anchors, descriptive alt text/captions, and transparent bylines.
5. **Freshness:** Update recency-sensitive pages (pricing, packaging, integrations, certifications) on predictable cadence with visible change logs.
6. **Provenance/UX hygiene:** Fast, accessible pages; clear authorship; durable URLs; minimal interstitial clutter.

Mini checklist (expandable into SOPs):

- Publish definitive facts and canonical “truth” tables.
- Ship honest buyer guides (comparisons, trade-offs) in buyer language.
- Add relevant structured data (product, FAQ, how-to); tag expert bylines.
- Align details with partner/marketplace pages; localize critical content.
- Remove rot and drift; enforce a single source of truth for specs.

6: Converting Attention to Attribution (Even when clicks are scarce)

Four complementary approaches

1. **Tagged answer journeys:** Where engines expose links, ensure tracking survives the handoff (UTM discipline; canonical deep links).
2. **Message-match fingerprinting:** When an answer rephrases a distinctive claim (“[Feature] cuts onboarding time by 42%”), detect subsequent on-site behavior that engages with the same claim to attribute influence without click.
3. **Geo/time experiments:** Nudge A-SOV in select markets; read downstream lift in branded search, direct landings, and CRM opportunities.
4. **Multi-touch modeling:** Assign partial credit to answer exposures alongside ads, email, and partner referrals; calibrate with experiment results.

Bottom line: Treat answer exposure as a top-funnel but high-intent touch. Build attribution that respects zero-click reality.

7: Case Snapshots (anonymized; methods abstracted)

Real outcomes with A-SOV/CQI as leading indicators. (Benchmarks available under NDA; collection methods withheld.)

B2B Cybersecurity (NA + EMEA):

- **Problem:** Low visibility in “best for mid-market X” prompts; inconsistent pricing facts across partner sites.
- **Interventions:** Authority coverage plan; pricing fact harmonization; proof-dense buyer guides.
- **Result (12 weeks):** A-SOV +19 pp (EN), +14 pp (DE); CQI from 48 → 67; Competitive Saliency shifted from “recommended second” to “peer” in 6 of 8 core prompts; qualified demos +11% QoQ (geo-matched).

Vertical SaaS (HRIS):

- **Problem:** Models summarized outdated packaging; negative sentiment around integrations.
- **Interventions:** Canonical truth tables; integration docs refresh; analyst coverage program; localized FAQs (ES/FR).
- **Result (10 weeks):** Narrative Alignment 2.3 → 4.1 (5-pt scale); A-SOV +13 pp; DI-D net uplift +9 pp vs. control; pipeline contribution +7% in treated markets.

Consumer Electronics (D2C):

- **Problem:** Product comparisons omitted brand; fragmented spec pages.
- **Interventions:** Specification consolidation; media review seeding; performance benchmarks with 3rd-party lab.
- **Result (8 weeks):** A-SOV +15 pp; CQI 52 → 71; retailer sell-through +6% in synchronized campaign window.

(Note: Snapshots illustrate the measurement model, not the proprietary steps taken to achieve these outcomes.)

8: Operating Model & Governance (How to run this quarter-to-quarter)

Roles & cadence

- **Owner:** Growth/SEO lead with Product Marketing and Comms.
- **Quarterly cycle:**
 - Re-sample prompts (by theme, locale, model)
 - Collect & score answers (A-SOV, CQI, Narrative)
 - Prioritize gaps (proof density, authority, fact drift)
 - Execute interventions (content, authority, alignment)
 - Run market experiments; refresh attribution; board-ready reporting

Guardrails

- **Legal/Comms:** Review claims, certifications, and third-party references; align with regulatory and IP policies.
- **Data & Ethics:** Respect robots/ToS; no scraping workarounds; document collection settings; red-team for prompt-injection and jailbreak risks.
- **Localization QA:** Native-speaker review for high-stakes pages; ensure fact parity across languages.

9: The Next 24 Months (Plan for change, not for a snapshot)

1. **More answers, fewer links.** Expect answer modules to expand and concentrate clicks into fewer, higher-authority citations; volatility will persist by vertical/intent.
2. **Agents compress the funnel.** Task and shopping agents will bundle consideration → decision in one thread; mention and citation become the gate to inclusion.
3. **Budget reallocation.** Dollars shift toward credibility assets (facts, authority, measurement) rather than pure content volume.
4. **Global by default.** Engines already answer in dozens of languages; teams that track only English lose edge (and revenue).
5. **Continuous tuning.** Model behavior changes; the brands that win will iterate scorecards and interventions every quarter.

10: What Jetstorm Provides (without exposing the “how”)

- **Answer Visibility Stack:** Instrumentation, dashboards, and scorecards for A-SOV, CQI, Competitive Salience, and Narrative Alignment by theme/model/locale.
- **Authority Orchestration:** A programmatic way to rebalance your Citation Authority Mix toward Tier-1/Tier-2 sources.
- **Truth Management:** Mechanisms to align and refresh verifiable facts across web properties, partners, and locales (pricing, SKUs, integrations, certifications).
- **Experiment & Attribution Framework:** Geo/time-boxed tests, message-match fingerprinting, and multi-touch models that connect answer presence to pipeline.
- **Compliance & QA:** Collection guardrails, accuracy workflows, and red-teaming for model-specific risks.

We share principles publicly and reserve “how we do it” for clients under NDA.

Appendix A — Metric Specifications

A-SOV (Answer Share of Voice)

- *Inputs:* Prompt list, model list, locale list.
- *Presence definition:* Named mention and/or link citation inside the synthesized answer (not only below-the-fold links).
- *Breakouts:* By model, locale, intent, funnel stage, and device.

CQI (Citation Quality Index)

- *Tiering policy (example):*
 - **Tier-1 (3):** Top analysts, standards bodies, Tier-1 media, your official docs, major developer portals.
 - **Tier-2 (2):** Recognized trade pubs, respected communities, high-quality comparison sites.
 - **Tier-3 (1):** Long-tail blogs, low-trust aggregators.
- *Normalization:* Divide by max possible score (total citations × 3).

Competitive Salience

- *Dimensions:* Co-mention frequency; relative framing (“best for...”, “budget option”, “enterprise leader”); evidence attached to each claim.
- *Alerting:* Flag if competitor is “recommended over” you in $\geq 20\%$ of high-intent prompts.

Narrative Alignment

- *Rubric signals:* ICP fit, value drivers, proof claims, pricing tier, deployment model, integration set, compliance status.
- *Output:* 0–5 score; explanations logged for remediation.

Appendix B — Experiment Design & Analytics

Design patterns

- **Matched market tests:** Pair regions with similar baseline demand; stagger interventions; use difference-in-differences (DiD) to estimate net effect.
- **Time-sliced tests:** Pre/post windows around discrete content/accountability interventions.
- **Sensitivity checks:** Exclude news spikes and product launches; validate with placebo prompts.

Key formulas

- **DiD effect:** $(\text{Test}_{\text{post}} - \text{Test}_{\text{pre}}) - (\text{Control}_{\text{post}} - \text{Control}_{\text{pre}})$ ($\text{Test}_{\text{post}} - \text{Test}_{\text{pre}} - (\text{Control}_{\text{post}} - \text{Control}_{\text{pre}})$).
- **Minimum detectable effect (MDE):** Use binary-proportion power analysis on A-SOV; calibrate sample sizes by $(\# \text{ prompts} \times \# \text{ models} \times \# \text{ locales} \times \# \text{ waves})$.
- **Attribution mixing:** Start with heuristic weights (exposure + branded search + direct + opportunity creation), then refit with experiment-observed lifts.

Dashboard suggestions

- **Tabs:** Overview | By Theme | By Model | By Locale | Authority | Experiments | Narrative QA.
- **Alerts:** A-SOV drop > 5 pp week-over-week in any high-intent theme; CQI Tier-1 share $< 30\%$ in any locale; new negative claim detected.

Appendix C — Operational SOPs (abbreviated)

Quarterly

- Refresh prompt set; add emerging tasks/use cases.
- Crawl & reconcile facts; rotate authority targets; re-certify high-impact pages.
- Run one geo-matched test per growth market; publish a board-ready Answer Scorecard.

Monthly

- Narrative QA on top themes; remediate inaccuracies.
- Expand localization coverage; validate parity with source facts.
- Update experiment notebooks and attribution weights.

Weekly

- Read A-SOV/CQI deltas; triage “answer gaps.”
- Inspect new citations and co-mentions; escalate misrepresentations to Comms/Legal.

Appendix D — Buyer's RFP Checklist (for GEO partners)

- Do you measure A-SOV and CQI by **model**, **theme**, and **locale**?
- Can you run **geo/time-boxed** tests and deliver statistically grounded readouts?
- How do you manage **fact parity** across sites, partners, and languages?
- What's your approach to **authority mix** (Tier-1/Tier-2) without spam or link schemes?
- How do you enforce **robots/ToS** compliance and mitigate model-specific risks?
- What are your **QA** and **red-team** processes for accuracy and prompt injection?

Closing

AI is the new infrastructure of the internet. It is rewriting the customer journey inside answers. With answer-first measurement (A-SOV, CQI), principled optimization (proof density, authority coverage, fact parity), and modern attribution (even without clicks), brands can convert answer presence into pipeline. Jetstorm's role is to make the invisible measurable—and then movable—without exposing the proprietary mechanics that drive the lift.

Sources for the evidence snapshot

- BrightEdge: AI Overviews frequency and trends; mid-teens prevalence after SGE testing.
- SE Ranking: AI Overviews research across 100k keywords; frequency dynamics 2024–2025.
- Search Engine Land: Zero-click behavior and share of clicks to the open web.
- Business of Apps: Perplexity usage and revenue run-rate.
- Gartner via VentureBeat: 2025 gen-AI spending forecast (~\$644B).
- IDC: AI market outlook to ~\$631B by 2028.
- iPullRank summarizing SE Ranking: Authority bias in AIO citations (top-10 correlation).
- Microsoft Copilot FAQ; Anthropic docs: Citation features and guidance.