

LLM Visibility Toolbox

Evidence-based tactics for getting cited in ChatGPT, Perplexity, Gemini, Claude, AI Overviews, and AI Mode.

By [Marcus Quinn](https://github.com/marcusquinn) (<https://github.com/marcusquinn>).

A Markdown-canonical playbook for AI search visibility. Use it to turn source evidence, page-type weighting, and answer-engine behaviour into roadmap-ready recommendations.

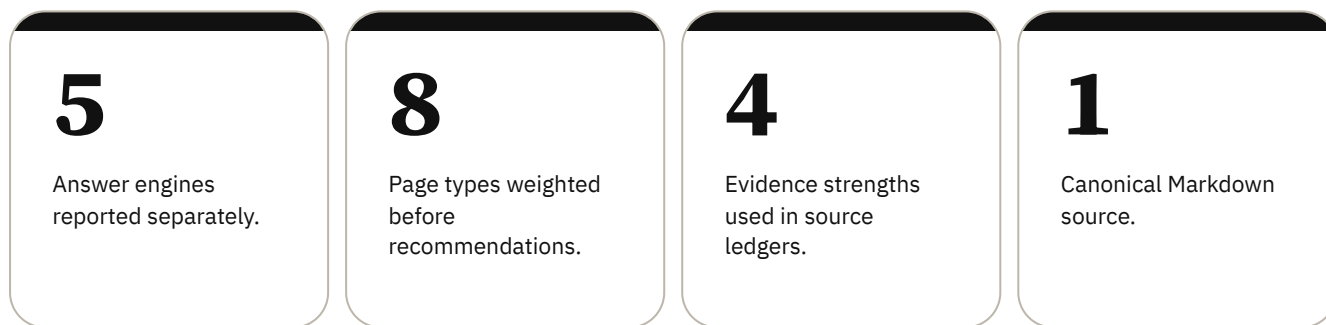
Audience: SEO, content, engineering, and leadership teams. Export rule: one HTML preview; PDF profiles for A4, Letter, and 16:9 decks.

Executive summary PEER-REVIEW

LLM visibility is an evidence system, not a single checklist. The best programmes make priority pages retrieval-ready, criteria-complete, citation-worthy, technically fetchable, and corroborated by third-party sources. **EVIDENCE:** VERIFIED

Recommendations must be weighted by page type. A homepage, SaaS feature page, comparison page, pricing page, article, product page, local/YMYL page, and research asset need different tactics, owners, and verification paths. **EVIDENCE:** VERIFIED

EVIDENCE: VERIFIED **EVIDENCE:** PARTIAL **EVIDENCE:** INFERRED **EVIDENCE:** MISSING



- PEER-REVIEW** Peer-reviewed or controlled comparison.
- STRONG** Large independent primary-data study.
- VENDOR** Vendor study with methodology and commercial incentives.
- PRACTITIONER** Practitioner evidence or field report.
- HYGIENE** Baseline technical implementation.

▲ **Changelog summary**

V4 adds schema-downgrade rationale, engine-specific reporting, source grouping, and diagram/equation fallbacks.

OPERATOR ACTION: collect source IDs first, then interpret findings into recommendations with owners, acceptance criteria, and rerun steps.

▲ **Action Prompt**

COPYABLE ACTION PROMPT



Reference this report action: collect source IDs first, then interpret findings into recommendations with owners, acceptance criteria, and rerun steps.

Guide me through the tools, resources, accounts, permissions, source material, and access needed to take this action. Break the work into numbered steps, call out any missing inputs before execution, include safe handling for credentials or confidential data, and finish with verification evidence I can capture.

Sources | Source ledger and evidence rules | appendices/source-ledger.md

Tactics | Highest-impact tactics and examples | weighted by page type

Matrix | Page-type matrix | required, conditional, avoid

Roadmap | Priority-card handoff | owner · effort · verification

Sources

Tactics

Matrix

Roadmap

Appendices

CONTENTS

LLM Visibility Toolbox

Executive summary

1. Highest-impact tactics

1.1 Weight before recommending

1.2 Critical

1.3 High

1.4 Medium

1.5 Low

2. Page-type matrix

2.1 Industry-fit reminder

3. Tactic card examples

3.1 Direct-answer opening

3.2 Impact

3.3 Evidence

3.4 Bot-friendly first fetch

3.5 Strong pattern

3.6 Weak pattern

4. Myths and caveats

4.1 Myth

4.2 Reality

5. Case studies

5.1 Industrial manufacturer

5.2 Healthcare comparison site

6. Roadmap template

6.1 Priority rule

6.2 Revenue page retrieval blocker

7. Verification checklist

8. Closing callouts

8.1 Combined finding

9. Sources

9.1 Prompt captures

9.2 Crawl evidence

9.3 Third-party profiles

9.4 Source-card rule

9.5 Ahrefs controlled schema study

9.6 Engine-overlap research

9.7 Buyer-research evidence

10. Appendices

1. Highest-impact tactics

1.1 Weight before recommending

Do not apply all tactics to all pages. Score page-type fit, retrieval eligibility, source proximity, corroboration, freshness, confidence, impact, and effort before roadmap sequencing.

1.2 Critical

Revenue page cannot be fetched or cited.

1.3 High

Claim lacks nearby evidence or source-card support.

1.4 Medium

Helpful page-type tactic with partial evidence.

1.5 Low

Hygiene, formatting, or monitoring improvement.

Tactic	Evidence	Best page types	Why it matters	Verification
Direct-answer opening	EVIDENCE: VERIFIED	Article, glossary, comparison, feature, local	Concise first-paragraph claims are easier to retrieve and cite.	Rendered first 300 words include answer, source, and updated date
Source cards near claims	EVIDENCE: VERIFIED	Research, comparison, YMYL, feature	Engines need nearby proof to trust and quote claims.	Source ID appears beside factual claim and in ledger

Tactic	Evidence	Best page types	Why it matters	Verification
Third-party corroboration	EVIDENCE: VERIFIED	SaaS, ecommerce, local, YMYL	Answer engines cross-check owned claims against outside sources.	Profile parity and source breadth review
Bot-friendly first fetch	EVIDENCE: VERIFIED	All priority pages	Hidden or blocked content cannot be cited.	Raw/rendered crawl, robots, sitemap, and logs
Entity consistency	EVIDENCE: PARTIAL	Homepage, about, local, profiles	Contradictory facts reduce answer confidence.	Canonical entity table and third-party parity
FAQPage schema	EVIDENCE: INFERRED	Hygiene only	Structured data helps clarity but does not replace visible evidence.	Schema validation plus visible-content check

2. Page-type matrix

Page type	Required tactics	Conditional tactics	Devalue or avoid
Homepage	Entity facts, category clarity, proof, crawlable nav	Original stats, comparison links	Long FAQ as primary GEO tactic
SaaS feature	Criteria block, use cases, integrations, proof	Demo video transcript, benchmark table	Generic benefit copy without source IDs
Pricing	Plan facts, constraints, comparison table	Purchase-relevant visible FAQ	Hidden pricing screenshots only
Comparison	Direct answer, feature/pricing table, alternatives, source cards	Third-party review quotes	Unsupported “best” claims
Article/guide	Direct answer, question headings, stats, expert quotes	Glossary sidebar, summary box	Thin filler or stale facts
Product/PDP	Specs, reviews, availability, canonical descriptions	Video transcript, product schema	Flat B2B SaaS checklist
Local/YMYL	Credentials, service area, policies, disclaimers	Practitioner bios, local citations	Unsupported advice
Research/report	Methodology, dataset, source cards, findings	Embeddable charts	PDF-only content without HTML summary

2.1 Industry-fit reminder

SaaS, ecommerce, local, and YMYL pages require different proof sources. Map the page type before assigning a tactic.

3. Tactic card examples

3.1 Direct-answer opening

- What: answer the query plainly in the first paragraph.
- Why: extractive systems need self-contained claims with nearby proof.
- How: pair answer, source ID, author/update date, and supporting table.
- Verify: rerun per-engine prompts and compare cited URL movement.

3.2 Impact

Direct-answer openings influence extraction quality, snippet usefulness, and the chance that a page is selected as a cited source.

3.3 Evidence

Verify with raw/rendered HTML, source-ID proximity, and per-engine prompt reruns.

3.4 Bot-friendly first fetch

- What: SSR or pre-render important content, allow relevant crawlers, and keep key text visible.
- Why: invisible or blocked content cannot be cited.
- How: compare raw HTML, rendered DOM, robots, sitemap, and logs.
- Verify: monthly crawl plus AI/search bot log review.

3.5 Strong pattern

Direct answer, evidence badge, source ID, visible methodology, updated date, and crawlable comparison table.

3.6 Weak pattern

Image-only proof, unsupported superlatives, client-rendered claims, and schema added without visible evidence.

4. Myths and caveats

4.1 Myth

Adding FAQPage schema is enough to become GEO-ready.

4.2 Reality

FAQPage is hygiene unless visible FAQ content genuinely fits page type and query fan-out.

TEXT CODE

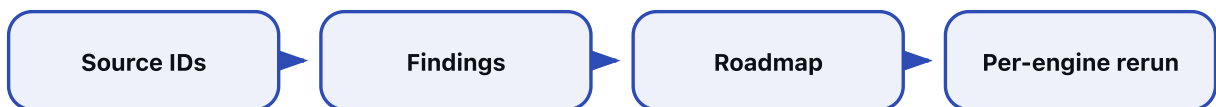


Worker brief: update /compare/example with source IDs S001-S004, visible comparison evidence, third-party corroboration, and retest steps. Acceptance: AIO, Gemini, ChatGPT, AI Mode, and Perplexity results are recorded separately.

MERMAID SOURCE FALLBACK



```
flowchart TD
  Sources[Source IDs] --> Findings
  Findings --> Roadmap
  Roadmap --> Rerun[Per-engine rerun]
```



Rendered Mermaid example, embedded as self-contained SVG.

Signal equation fallback: $AI \setminus \text{visibility} = \text{retrieval} + \text{evidence} + \text{corroboration}$.

TEXT CODE



Written by Dr. Jane Doe, PhD
Principal Data Scientist, ExampleCo

Use this block for named experts, source credentials, and profile links.

On-page evidence — 72%

Technical retrieval — 64%

Authority corroboration
— 58%

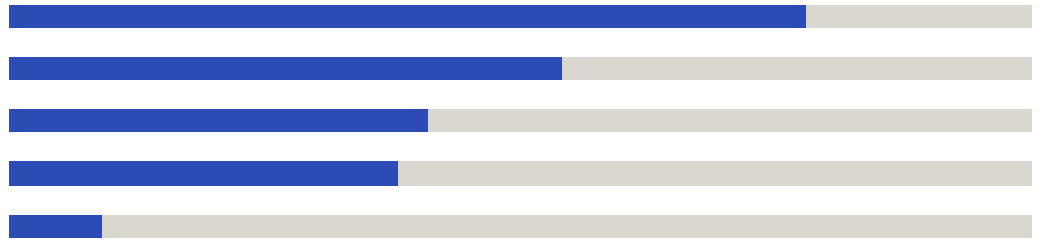
AI Overviews — 78%

Gemini — 54%

ChatGPT — 41%

AI Mode — 38%

Perplexity — 9%



Strong reports separate observed facts from interpretation, then turn only verified or clearly labelled partial evidence into roadmap items.

5. Case studies

5.1 Industrial manufacturer

Result: measurable AI referral growth after direct-answer restructuring and third-party corroboration.

Tactics applied: original benchmarks, source-card evidence, bot-friendly rendering, and trade-publication mentions.

5.2 Healthcare comparison site

Result: cited answers appeared across multiple answer engines after visible expertise and profile parity fixes.

Tactics applied: practitioner bylines, review methodology, source-backed tables, and monthly prompt reruns.

6. Roadmap template

6.1 Priority rule

Start with revenue pages that fail retrieval eligibility or evidence proximity before optional schema enhancements.

6.2 Revenue page retrieval blocker

Use a priority card when a recommendation must carry priority, owner, due date, source IDs, and verification in one executive-scannable block. Pair with the source ledger and a worker-ready implementation brief.

Priority	Recommendation	Applies to	Owner	Verification	Source IDs
P0	Fix retrieval blockers on revenue pages.	Homepage, pricing, feature, PDP, local	SEO + engineering	Raw/rendered crawl, robots, sitemap, logs	S002, S005
P1	Add source cards and original evidence.	Comparison, article, research/report	Content + subject expert	Source ledger and citation checks	S001, S003
P1	Build third-party corroboration.	SaaS, local, ecommerce	Marketing/PR	Profile parity and source breadth	S004
P2	Improve schema and metadata.	All page types	SEO + engineering	Schema validation plus visible-content check	S002

7. Verification checklist

- Validate evidence badges and source IDs before export.
- Render HTML with the chosen DESIGN.md-backed template.
- Review table wrapping, badge visibility, source-card readability, and sticky TOC behaviour.
- Export A4/Letter PDF for documents and 16:9 PDF for decks.
- For client-custom reports, rerun live evidence collection before interpretation.
- For recurring reports, create a custom routine with deterministic collection and report-agent interpretation.

8. Closing callouts

8.1 Combined finding

AI visibility reporting should end with the fewest useful recommendations: retrieval blockers, evidence proximity, third-party corroboration, and monitoring. Keep panels for important emphasis; use plain bullets and tables for normal content.

9. Sources

Primary sources

9.1 Prompt captures [↗](#)

AIO, Gemini, ChatGPT, AI Mode, and Perplexity prompt evidence stored separately.

9.2 Crawl evidence [↗](#)

Raw/rendered crawl export with retrieval eligibility notes.

Corroboration sources

9.3 Third-party profiles [↗](#)

Review, directory, community, partner, and media source parity checks.

Source ID	Evidence type	Use in report	Verification
S001	Prompt capture	Per-engine citation presence	AIO, Gemini, ChatGPT, AI Mode, and Perplexity recorded separately
S002	Raw/rendered crawl	Retrieval eligibility	Important claims visible on first fetch
S003	Page inventory	Page-type weighting	URL mapped to homepage, feature, comparison, article, local, PDP, or report
S004	Third-party profile	Corroboration strength	Facts match owned canonical entity table

Source ID	Evidence type	Use in report	Verification
S005	Analytics/search data	Business value and priority	Priority URL cluster tied to demand or revenue

9.4 Source-card rule

Every roadmap item should cite source IDs, observed date, confidence, owner, and the command or routine that verifies completion.

PUBLIC ARTIFACT RULE

Do not export private URLs, raw transcripts, screenshots, local paths, or client names. Publish source IDs and redacted summaries; keep raw evidence in approved secure storage.

9.5 Ahrefs controlled schema study [↗](#)

Schema is treated as technical hygiene rather than a primary AI-visibility growth lever.

9.6 Engine-overlap research [↗](#)

Low overlap between AIO, Gemini, ChatGPT, AI Mode, and Perplexity requires per-engine reporting.

9.7 Buyer-research evidence [↗](#)

Answer engines increasingly influence discovery and shortlisting, so reports separate visibility from conversion value.

▲ **How source IDs become recommendations**

1. Capture source evidence before writing findings.
2. Map each source to supported claims and page types.
3. Score recommendations by impact, confidence, effort, and verification path.
4. Include unresolved gaps in the appendix rather than presenting them as facts.

10. Appendices

[Source ledger appendix](#) MD

[Prompt set appendix](#) MD

[Client audit example](#) HTML

[Style previews](#) HTML

V4 · COMPILED MAY 2026 FROM SOURCE-LED EVIDENCE · INTERNAL TOOLKIT

© 2025-2026 Marcus Quinn. Licensed under MIT.