

Report Style Showcase

Evidence-based report components for AI-search, SEO/GEO, and operational audit exports.

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

Component stress-test for DESIGN.md-backed report styles. Render this same Markdown through different templates to compare typography, palette, spacing, borders, cards, tables, badges, code blocks, and print profiles.

34

Renderer templates.

16:9

Slide PDF profile.

A4

Document PDF profile.

MD

Canonical source.

Executive summary

This showcase demonstrates the report components expected in the LLM Visibility Toolbox examples: a non-numbered executive summary, numbered H2 chapters, plain H3 subheadings, sources grouped by type, badge keys, checklists, code/block templates, callouts, accordions, and a version summary.

PEER-REVIEW

Peer-reviewed paper or controlled comparison with documented methodology.

STRONG

Large primary-data analysis from an independent source.

VENDOR

Vendor-published study with methodology, but commercial conflict of interest.

PRACTITIONER

Practitioner report or anecdote, often without a control group.

HYGIENE

Technical baseline that is not experimentally measured but follows from how the bots work.

▲ What changed in v4

Schema was downgraded after controlled evidence; engine divergence is now structural; source sections are grouped by evidence type; reports keep Mermaid and LaTeX as readable fallbacks unless local renderers pre-render them.

ACTION: review the numbered chapters below and compare the generated HTML against the source Toolbox patterns.

▲ Action Prompt

COPYABLE ACTION PROMPT



Reference this report action: review the numbered chapters below and compare the generated HTML against the source Toolbox patterns.

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.

CONTENTS

Report Style Showcase

Executive summary

1. Why this matters now

2. Highest impact, most validated

2.1 Earned media on third-party platforms

3. On-page content tactics

4. Technical tactics

4.1 Info panel

4.2 Impact panel

4.3 Action panel

4.4 Tactic card

4.5 Good row

4.6 Bad row

4.7 Myth

4.8 Fact

5. Off-page and authority

6. Format and experimental tactics

6.1 SEO myths called out

7. Measurement and reporting

7.1 High-priority visual checks

7.2 Refresh checklist per page

8. What does not work

8.1 Combined finding

9. Productizing for clients

10. Case studies

10.1 Industrial manufacturer

10.2 Healthcare comparison site

11. Sources

11.1 Source A

11.2 Source B

11.3 Source C

11.4 GEO: Generative Engine Optimization

11.5 News source citing patterns in AI search systems

11.6 Ahrefs: schema markup has no impact on AI visibility

11.7 Growth memo: the consensus gap

11.8 G2: the answer economy research

12. Appendices

CHAPTER 1

1. Why this matters now

Use [anchor links](#) ([#priority-and-checklist](#)), [appendix links](#) ([../llm-visibility-toolbox/report.html](#)), numbered steps, accordions, coloured panels, and source cards in the same canonical Markdown.

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

Evidence values should read as plain **Evidence:** text followed by a colour-coded mini badge for the value only.

Table stress

Cards

Priority

2. Highest impact, most validated

Use plain narrative and bullets when that is clearer than a panel. Panels are reserved for warnings, action blocks, source cards, or high-emphasis evidence.

2.1 Earned media on third-party platforms **STRONG**

Six converging studies point to third-party mentions as a stronger AI visibility signal than isolated owned-page edits.

- Build a quarterly cadence across trade press, community, video, and partner surfaces.
- Pair each campaign with source IDs and prompt reruns.
- Verify movement separately in AIO, Gemini, ChatGPT, AI Mode, and Perplexity.

ACTION: coordinate one trade article, one community thread, one video transcript, and one partner citation within the same quarter.

▲ Action Prompt

COPYABLE ACTION PROMPT



Reference this report action: coordinate one trade article, one community thread, one video transcript, and one partner citation within the same quarter.


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.

A plain bullet section should remain plain:

- Robots.txt and crawlability check for key AI bots.
- Render test with a representative answer-engine user agent.

- Schema audit on priority pages.
- Prompt list from Search Console, support tickets, and customer interviews.
- Baseline share of voice, citation rate, and sentiment per engine.

3. On-page content tactics

Component	Stress condition	Expected result
Evidence badge	Long table cell with badge EVIDENCE: 	Badge stays readable and does not split words.
Facts table	Multiple columns with prose	Table remains usable in HTML and constrained in print.
Source card	Evidence note near claim	Card is visually distinct from normal paragraphs.
Sidebar	Many headings	Sticky TOC remains secondary to content and active link updates.

4. Technical tactics

4.1 Info panel

Use info panels for caveats, assumptions, and reading guidance that should not become recommendations.

4.2 Impact panel

Why this changes the plan: a high-impact finding changes sequencing, budget, owner, or acceptance criteria.

4.3 Action panel

Next action: turn the finding into a concrete implementation step with owner, proof path, and rerun command.

▲ Action Prompt

COPYABLE ACTION PROMPT



Reference this report action: turn the finding into a concrete implementation step with owner, proof path, and rerun command.

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.

4.4 Tactic card

- What: compact recommendation summary.
- Why: links action to evidence.
- How: gives implementation shape.
- Verify: names the acceptance check.

4.5 Good row

Crawlable claims, source IDs, direct answers, and page-type weighting.

4.6 Bad row

Unsupported claims, hidden content, generic tactics, and missing verification.

4.7 Myth

Every page needs the same GEO checklist.

4.8 Fact

Page type determines which tactics are useful,

conditional, or noise.

▲ **How to read this document (evidence badges)**

Every tactic carries a badge. Use RCT/academic for controlled research, strong primary data for large independent data, vendor study where methodology exists but incentives are commercial, practitioner for field evidence, and hygiene for baseline technical work.

▲ **Key methodology caveat**

A visible uplift in one engine is not proof of universal AI visibility. Keep AIO, Gemini, ChatGPT, AI Mode, and Perplexity separate until the closing synthesis.

Quotes highlight expert evidence, user language, or a decision constraint without turning it into a recommendation.

Author block template

TEXT CODE



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

Jane led the data platform team from 2019 to 2024 and now researches LLM retrieval.

LinkedIn | Google Scholar | Personal site

Render command

TEXT CODE



```
.agents/scripts/report-render-helper.sh render report.md --template lottiefiles --pdf-profile slides-16-9-2 --output report.html
```

Mermaid fallback

MERMAID SOURCE FALLBACK



```
flowchart LR
  SourceLedger --> Finding
  Finding --> Recommendation
  Recommendation --> Verification
```



Rendered Mermaid example, embedded as self-contained SVG.

LaTeX fallback

LATEX SOURCE FALLBACK



$\text{LLM visibility} = \alpha \text{ citations} + \beta \text{ mentions} + \gamma \text{ retrieval} - \delta \text{ decay}$

$\text{LLM visibility} = \alpha \text{ citations} + \beta \text{ mentions} + \gamma \text{ retrieval} - \delta \text{ decay}$

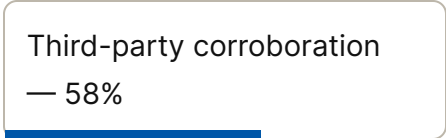
Rendered LaTeX example, embedded as self-contained HTML.

Inline LaTeX fallback: $\text{visibility} = \alpha \text{ citations} + \beta \text{ mentions} + \gamma \text{ retrieval}$.

Citation readiness — 72%



Third-party corroboration
— 58%



Retrieval eligibility — 81%



CHAPTER 5

5. Off-page and authority

Authority work belongs outside the site as much as on it: profile parity, trusted third-party mentions, practitioner credentials, and community proof all support retrieval and citation decisions.

6. Format and experimental tactics

6.1 SEO myths called out

Claims that circulate widely but cannot be traced to a primary source, or are contradicted by controlled evidence, should be called out explicitly.

“Schema markup alone creates citation uplift.” Contradicted by controlled or near-controlled studies; ship schema as hygiene.

“Longer content always gets cited more.” Engine-dependent; content depth helps only when it improves answer density and source usefulness.

7. Measurement and reporting

7.1 High-priority visual checks

Review spacing, table width, no-wrap badges, active TOC highlighting, print CSS, and component contrast.

7.2 Refresh checklist per page

Replace at least one statistic with newer data.

Add or revise one example.

Remove outdated tool or vendor references.

Update screenshots if UI has changed.

Update `dateModified` in schema.

Add visible “Last updated” line.

8. What does not work

8.1 Combined finding

AI is now a discovery layer, but engines disagree on sources. Tracking only mentions or only citations misses the retrieval gap. Keep the final synthesis short, source-backed, and tied to the next action.

CHAPTER 9

9. Productizing for clients

Use the same Markdown-first report structure for one-off audits, monthly retainers, lead magnets, and routine handoffs. Keep deterministic evidence collection separate from interpretation.

10. Case studies

10.1 Industrial manufacturer

Result: monthly AI referral traffic grew from near-zero to a measurable assisted-conversion channel.

Tactics applied: direct-answer page restructure, original technical benchmarks, schema hygiene, and trade-publication mentions.

10.2 Healthcare comparison site

Result: citations appeared across Google AIO, ChatGPT, and Gemini after entity facts and expert review were made visible.

Tactics applied: YMYL author bylines, source-backed comparison tables, third-party profile parity, and prompt reruns.

11. Sources

Primary evidence

11.1 Source A

Prompt capture, crawl export, and source ledger row.

11.2 Source B

Third-party corroboration and profile parity note.

Supplementary evidence

11.3 Source C

Appendix file, screenshot reference, or companion report.

11.4 GEO: Generative Engine Optimization [↗](#)

The first peer-reviewed baseline. Tactics tested on 10k queries.

11.5 News source citing patterns in AI search systems [↗](#)

Top source concentration and citation dynamics.

11.6 Ahrefs: schema markup has no impact on AI visibility [↗](#)

1,885 vs 4,000 controls, difference-in-differences. Source used to downgrade schema from growth lever to hygiene.

11.7 Growth memo: the consensus gap [↗](#)

Only a small share of cited URLs overlap across engines; engine-specific reporting is required.

11.8 G2: the answer economy research [↗](#)

B2B buyers increasingly start with answer engines, so reports separate discovery, shortlist, and conversion evidence.

12. Appendices

 [Source ledger appendix](#) **MD**

 [Client audit example](#) **HTML**

 [Example browser](#) **HTML**

V4 · COMPILED MAY 2026 FROM 70+ PRIMARY SOURCES · INTERNAL TOOLKIT

© 2025-2026 Marcus Quinn. Licensed under MIT.