

CONTENT AUDIT



THE BRIEF

A global automotive client asked SJR to test how well its newsroom content would be found, cited, and relied on by AI-driven answer engines. The goal: identify content gaps, technical bottlenecks, and quick wins that would improve discoverability and reduce accuracy risk.



WHAT WE DID

We built a custom knowledge base and ran a two-part audit

- ✓ A content evaluation across over 500 core questions
- ✓ A technical AI-Search Optimization assessment after scraping ~10,400 public pages.

We also benchmarked public models (e.g., GPT-4) against the client's ground truths to surface hallucination risk and attribution gaps.



KEY FINDINGS

Content correctness

- **Answers:**
68% matched ground truth
16% incorrect
16% no answer returned
- **Biggest issue:**
Content age (older pages, product pages)
- **Largest subject gaps:**
Corporate, innovation

Technical readiness

- **AISO readiness:**
Low (4.2/10)
- **Issues:**
Link density, server fetch times, media transcripts/alt text, incomplete JSON-LD



DELIVERABLES

Scored dashboard

Overall readiness + sub-scores

An executive one-pager

Highlighting the top quick wins.

Prioritized roadmap

Quick wins (FAQ blocks, answer-first summaries, JSON-LD snippets), mid-term content refreshes (canonicalization, taxonomy), and technical fixes (semantic HTML, TTFB/CDN, media captions).

Public-AI comparison outputs

Reveal where external models fill gaps with unofficial sources (hallucination risk).

Measurement playbook

Recommend randomized page-level holdouts and KPI definitions (citation share, referral CTR, hallucination/error rate).



WHY IT MATTERED

The audit turned an abstract risk (AI misattribution and hallucination) into a prioritized action plan: refresh the highest-impact pages, fix a small set of technical blockers, and instrument tests that prove lifts in visibility and conversion. That roadmap made it clear where to invest first to protect brand trust and capture higher-quality referrals.