My First AI Visibility Audit & Report Planner

JvGLabs |

Helping agencies and their clients become discoverable and recommendable across AI search ecosystems — providing the infrastructure, automation, and expertise to navigate the shift from traditional SEO to AI Visibility.

A Step-by-Step Working Document for Agency Strategists

Who this is for: Agency strategists, SEO leads, and account managers who want to move beyond keyword-based reporting and offer a new, high-value service: the AI Visibility Audit.

What problem it solves: You have read about entities, knowledge graphs, and LLM citations. You understand *why* it matters. But when you sit down to actually build your first audit for a real client, the gap between theory and execution feels enormous. This planner closes that gap.

How to use it: Pick one client. Work through each phase in order, filling in every field. By the end, you will have a complete strategic brief — the foundation for a white-label AI Visibility Report you can deliver and sell.

What you get: A focused, scoped, and narrative-driven action plan for your first pilot audit — plus a repeatable framework you can use for every client after that.

Why this connects to AI Visibility: Search and AI systems increasingly rank and recommend content based on entities (the specific people, concepts, products, and organizations a brand is associated with) and their relationships — not just keywords. Building a brand's semantic footprint (the web of entities and topics AI systems associate with that brand) is the core transformation that determines whether a brand gets cited, recommended, or ignored by AI. This planner gives you the concrete starting point for that work.

Quick Win: Client Selection Scorecard

Not every client is the right candidate for a pilot AI Visibility Audit. Starting with the wrong client leads to wasted effort and underwhelming results. Use this scorecard to choose wisely. The ideal client scores at least 3 out of 5.

Selection Criterion	Check if True	Why It Matters
Their leadership is actively asking about AI's impact on their business.	☐	Signals executive buy-in — you will not have to sell the *concept* before selling the *service*.
They operate in a competitive B2B or high-consideration B2C market where expertise matters.	☐	These markets depend on authority and trust — exactly what entity-based visibility builds.
They have a substantial library of existing content (blog posts, whitepapers, case studies).	☐	More content means more data to audit and more opportunities to find quick wins.
They understand the difference between traffic and authority.	☐	They will value being *recommended* by an AI, not just ranking for a keyword.
They have expressed frustration with the limitations of keyword-based reporting.	☐	They are already looking for something better — you are about to show them what that looks like.

Score: _____ / 5

Selected Client for Pilot Audit: _____

Why this client (one sentence): _____

Outcome: You have chosen a client who is both ready and likely to appreciate the value of what you are building.

Phase 1: LLM Benchmark Scoping

Goal: Define the exact parameters for measuring your client's current visibility inside conversational AI systems. Without clear scope, you will drown in data. This phase prevents that.

1.1 Client & Competitor Domains

Role	Domain	Notes
Your Client	_____	
Primary Competitor	_____	The one your client loses deals to most often.
Secondary Competitor	_____	An aspirational competitor or emerging threat.

1.2 Core Topics to Query (Non-Branded)

These should be problem-oriented or solution-oriented questions that a potential *customer* of your client would ask an AI assistant. Do not use branded terms — the point is to see whether AI recommends your client organically.

#	Query / Topic	Why This Matters for the Client
1	_____	_____
2	_____	_____
3	_____	_____
4	_____	_____
5	_____	_____

Tip: Start with five queries. You can expand later. Five well-chosen queries will reveal more than fifty vague ones.

1.3 Target LLMs for Testing

Check the platforms you will query. Testing across multiple LLMs (large language models — the AI systems behind tools like ChatGPT, Gemini, and Perplexity) gives a more complete picture, since each model draws on different training data and retrieval methods.

- ☐ ChatGPT (GPT-4 or later)
- ☐ Perplexity
- ☐ Google Gemini / AI Overviews (Google's AI-generated search summaries)
- ☐ Claude
- ☐ Other: _____

1.4 Important Limitation to Acknowledge

LLM outputs are non-deterministic (the same question can produce different answers each time). A single query is a snapshot, not a statistic. For meaningful benchmarking, plan to run each query multiple times across different sessions and record whether your client is mentioned, how they are described, and in what position relative to competitors. Document the date and model version for each query — results will shift as models are updated.

Phase 1 Outcome: You now have a precise, bounded list of queries, competitors, and platforms. This becomes your benchmark protocol.

Download Free Planner

Phase 2: Entity Gap Analysis Scoping

Goal: Identify the core semantic territory where your client needs to build authority, and define the battleground for comparison against competitors.

2.1 Client's Core "Expertise" Entity

An entity (in this context, a specific concept, topic, person, product, or organization that AI systems recognize and use to organize knowledge) is the building block of AI understanding. What is the one primary concept your client needs to be definitively associated with?

Core Entity: _____
(Example: "Enterprise Cybersecurity Solutions," "Sustainable Supply Chain Management," "Pediatric Telehealth")

2.2 Key Thematic Areas to Compare

These are the sub-topics that make up your client's core expertise. Think of them as branches of a knowledge graph (a structured map of entities and how they relate to each other) that AI systems use to decide who is an authority on what.

#	Thematic Area	Your Client Has Content?	Competitor Likely Covers This?
1	_____	☐ Yes ☐ No ☐ Partial	☐ Yes ☐ No ☐ Unknown
2	_____	☐ Yes ☐ No ☐ Partial	☐ Yes ☐ No ☐ Unknown
3	_____	☐ Yes ☐ No ☐ Partial	☐ Yes ☐ No ☐ Unknown

2.3 Entity Gap Hypothesis

Before running any analysis, write down your gut feeling. Having a hypothesis makes the data more meaningful — you will either confirm it or discover something surprising.

"I expect the biggest entity gap will be around: _____

because: _____"

Phase 2 Outcome: You have defined the scope for your knowledge graph comparison. This prevents analysis paralysis (the tendency to keep researching instead of acting) and ensures your audit focuses on the most strategically valuable territory.

Phase 3: Semantic Content Audit Scoping

Goal: Select a small, representative sample of your client's most important content and define what you will measure. A semantic content audit (evaluating how well content communicates meaning to machines, not just humans) at scale starts with a focused pilot.

3.1 Select 5–10 Key Content Assets

Choose pages that *should* be driving revenue or establishing authority. These are the pages where poor machine understanding costs the most.

#	URL of Content Asset	Topical Depth	Entity Coverage	Schema Markup
1	_____	☐	☐	☐
2	_____	☐	☐	☐
3	_____	☐	☐	☐
4	_____	☐	☐	☐
5	_____	☐	☐	☐

What these columns mean:

- **Topical Depth** — Does the content thoroughly cover its subject, including related sub-topics, or does it only skim the surface?
- **Entity Coverage** — Does the content mention and explain the key entities (specific concepts, products, standards, people) that an AI system would expect to find on a page about this topic?
- **Schema Markup** (structured data embedded in a page's code that explicitly tells search engines what the content is about) — Does the page use JSON-LD (a lightweight format for structured data) or other schema to make its entities machine-readable?

3.2 Define Your Scoring Approach

For each dimension above, decide in advance what "good" and "poor" look like for this specific client and industry. This does not need to be a complex formula — a simple three-level scale works well for a pilot.

Dimension	Score: Strong	Score: Adequate	Score: Weak
Topical Depth			
Entity Coverage			
Schema Markup			

Important: Any scoring framework you build here is a simplified heuristic (a practical rule of thumb, not a precise measurement). Real-world scores are affected by factors like content freshness, how frequently AI training data is updated, the specific NLP libraries (software tools for processing human language) you use, and the evolving criteria AI systems apply. State these limitations clearly in your client report. A directionally correct audit presented honestly is far more valuable than false precision.

Phase 3 Outcome: You have a targeted, manageable list of content to analyze — and a clear rubric for scoring it. This makes the audit actionable instead of overwhelming.

Real-Life Scenario: Avoiding the "Data Dump"

The Pitfall Most Agencies Fall Into

You run sophisticated analyses. You extract entity lists, chart LLM mention frequencies, score content for semantic depth. You deliver a 30-page report packed with charts. The client is briefly impressed — then confused. They do not know what to do next. The report sits in a folder. The retainer is not renewed.

This is the "data dump" trap.

The Pro Move: Lead With "So What?"

For every data point in your report, answer three questions before including it:

1. **What does this mean in plain language?**
2. **Why should the client care?**
3. **What should we do about it?**

Example — Weak vs. Strong Framing:

Weak (Data Dump)	Strong (Strategic Narrative)
"Here is a list of 50 entities associated with your competitors that are absent from your content."	"Competitors own the conversation around 'supply chain resilience.' This entity gap is a key reason AI systems recommend them for enterprise logistics queries, not you. Our first priority: build authority around this topic, starting with these three existing pages that can be upgraded quickly."
"Your content scored 4/10 on topical depth for 'cloud migration.'"	"Your 'cloud migration' page covers the basics but misses the sub-topics AI systems expect — specifically hybrid cloud orchestration and data sovereignty. Adding these sections will make this page substantially more useful to both human readers and AI."

The Takeaway: The report that renews the contract is not the one with the most data. It is the one that tells a clear story: *here is where you stand, here is where the opportunity is, here is what we do first.*

Phase 4: Building Your Narrative

Goal: Before you touch any data, define the story you expect to tell. This keeps your entire audit oriented toward a strategic conclusion, not just a collection of findings.

4.1 The Key Story This Data Will Tell

(Write one sentence. Example: "The client is perceived as a niche player in one region, but the data will reveal an opportunity to become a recognized authority across a broader market.")

_____

_____

4.2 The Biggest Opportunity You Expect to Uncover

(Write one sentence. Example: "No competitor has built entity authority around 'AI-powered compliance' — the client can own this emerging space.")

_____

_____

4.3 The Immediate Next Step for the Client

(Write one concrete action. Example: "Optimize the three service pages identified in the content audit to close the most critical entity gap within 30 days.")

_____

_____

4.4 Narrative Reality Check

Your hypothesis may be wrong. That is fine — and actually valuable. Use the checkboxes below after completing your data gathering to validate or revise your narrative.

- ☐ The data confirmed my expected story.
- ☐ The data partially supported it — I adjusted the narrative as follows:
  _____
- ☐ The data told a completely different story — and it is actually more compelling:
  _____

Phase 4 Outcome: You have a working thesis for your report. Combined with the scoping from Phases 1–3, you now have a complete strategic brief for your first AI Visibility Audit.

Why This Matters for Your Agency's AI Visibility Service

The shift from traditional search to AI-driven discovery is not a distant prediction — it is happening now. According to industry surveys (cited via Perplexity AI aggregation), approximately 91% of agencies report using AI in some form, but only around 39% have deeply integrated it into their workflows. Entity-rich content has been observed to earn substantially more AI citations than keyword-optimized content alone (typical estimates suggest approximately 3x more, though exact figures vary by industry and model).

This means there is a window of opportunity. Most agencies are still selling keyword rankings. The agency that can walk a client through a rigorous AI Visibility Audit — showing them exactly where they are invisible to AI, exactly where the gaps are, and exactly what to do about it — will own the most valuable positioning in the market.

This planner is your starting point for building that service.

Your Blueprint Is Complete — Here Is What to Do Next

Summary: You have scoped your first AI Visibility Audit across all four phases: LLM benchmarking, entity gap analysis, semantic content audit, and narrative strategy. You have a selected client, defined queries, identified competitors, chosen content assets, and drafted your story.

Next Step 1: Start Building With the Right Tools

Begin the data gathering for Phase 1 by running your benchmark queries using the LLMs you selected. For a full technical walkthrough — including workflows, templates, and the complete four-part methodology — explore the AI Visibility Audit framework on JvGLabs.

→ **Access the full AI Visibility Audit framework:** [jvglabs.com/ai-visibility-audit-framework]

Next Step 2: Get Expert Guidance for Your Pilot

If you want hands-on support scoping your first audit, reviewing your findings, or building a white-label reporting workflow for your agency, connect with an AI visibility strategist who specializes in entity-based optimization and LLM discoverability.

Save or print this planner. Use it as your working document during the audit — and as a repeatable template for every new client engagement.

JvGLabs |
Helping agencies and their clients become discoverable and recommendable across AI search ecosystems — providing the infrastructure, automation, and expertise to navigate the shift from traditional SEO to AI Visibility.

JvGLabs — AI visibility architecture
Designing the structural and semantic foundations that make brands findable, citable, and recommendable across AI-driven search and discovery systems.

Download Free Planner