

WHAT EACH AUDIT TIER DELIVERS

Four audits, four niches — see what each tier adds.

One example per tier, each for a different fictional brand in a distinct category, so you can see exactly how depth escalates from a free snapshot to a comprehensive plan. Every figure is sourced and method-noted.

 **Illustrative samples** — not real client engagements. Figures are representative; method noted on each.

Google AI surfaces — measured separately (floor, not a ceiling)

We report Google AI Overviews as its own surface, never blended into a cross-engine score — it only fires on some searches, and a click lands in your analytics as ordinary organic traffic, so we treat measured AI referral as a floor, not a ceiling.


FREE TEARDOWN

\$0

Brand A · observability · "Lumen Observe"

The one-page snapshot

SHARE OF MODEL · "BEST OBSERVABILITY TOOL FOR KUBERNETES" · ADAPTIVE PER-ENGINE SAMPLING



5th of 5. Source: illustrative teardown · 5 prompts × 3 engines · Jun 2026.

Three quick wins

01 Publish a "high-cardinality cost" comparison page
Engines cite a cost-comparison source you're absent from.

02 Add Organization + Product schema
Give engines clean structured data to parse.

03 Answer-first intro on your top docs page
Lead with the recommendation engines can lift.

STARTER \$500


Brand B · vector database · "Vectyr DB"

+ Adds over Free: 10 prompts · 2-3 engines · AI-readiness scorecard · top-10 fixes

Technical AI-readiness scorecard



AI Visibility Score 33 / 100 ("Building"). Source: illustrative Starter audit · 10 prompts × 3 engines · 95% CI · Jun 2026.



Top-10 fixes (excerpt)

- 01 Publish a latency/recall benchmark page


Owned source engines can cite for "production RAG".
- 02 Earn a GitHub README comparison mention

Engines pull from repo docs in 6/10 runs.
- 03 + 7 more, prioritised by effort × impact

Full list in the deliverable.

Competitor share-of-model leaderboard

● MEASURED · 24 PROMPTS × 5 ENGINES



4th of 4 · widest gap on Grok. Source: illustrative Full audit · 24 prompts × 5 engines · 95% CI · Jun 2026.

SEO foundation findings

- 14 keywords in top 10; thin comparison coverage
- Core Web Vitals pass; schema missing on key pages
- Backlink gap vs Pistons on benchmark content

30/60/90 roadmap

- 0–30: benchmark page + schema
- 31–60: 3 listicle placements + 4 articles
- 61–90: digital PR + re-measure

COMPREHENSIVE


\$4,500

Brand D · feature flags · "Flagstone"

+ Adds over Full: top-5 page rewrites · schema specs · off-site placement plan · directory submission list

Everything in Full, plus the build plan

SHARE OF MODEL · "BEST FEATURE-FLAG TOOL FOR SMALL TEAMS"



Top-5 page rewrites + schema specs

- Homepage, /pricing, 3 docs — answer-first rewrites
- Article + FAQ + SoftwareApplication JSON-LD specs
- Entity / sameAs map for the brand

Off-site + directory plan

- Placement plan: G2, Capterra, comparison listicles
- Directory submission list (30+ curated)
- Digital-PR angle list with target outlets

DELIVERED A ready-to-execute build plan

Comprehensive is the only tier that ships the actual rewrites, schema, and submission lists — not just findings.

What each tier adds, at a glance

TIER	PROMPTS × ENGINES	HEADLINE DELIVERABLE
Free	5 × 3	Snapshot + 3 quick wins
Starter	10 × 2-3	Readiness scorecard + top-10 fixes
Full	20-30 × 5	Leaderboard + SEO + 30/60/90
Comprehensive	20-30 × 5	Full + rewrites + schema + off-site plan

See your own numbers — free.

Start with a free teardown; the audit fee is credited to your first retainer.

Get a free teardown

Clear Cited · all four brands are fictional samples · figures representative, not measured from real clients · method: adaptive per-engine sampling, metric-matched 95% CI (presence Wilson / share bootstrap).