



MONTH-ONE SEO AUDIT

SEO Audit & 90-Day Game Plan

A full audit of the practice's organic search footprint — where it stands, what's quietly costing it its most valuable searches, and the 90-day plan to fix it.

PREPARED FOR

A multi-location specialty practice

PREPARED BY

Kava Digital Marketing

SCOPE

Audit + 90-day plan

SCOPE & DATA SOURCES

What we audited

A full audit of the practice's organic search presence: the keyword footprint and where the money intent actually sits, the technical health of the site, the on-page and content architecture that decides which pages rank, the competitive gap, and the 90-day plan that comes out of it. Every figure carries the source it came from and the comparison window it was measured over.

Sources across this document: Google Search Console (GSC) · SEMrush (Domain Overview, Organic Research, Backlink Analytics) · Google Lighthouse / Chrome UX Report (CWV) · live on-page DOM crawl · sitemap crawl · Google Business Profile. Baselines pulled over a single audit cycle; comparison window noted per figure.

Executive summary

On paper, this practice looked healthy: tens of thousands of ranking keywords, a six-figure quarterly click count, and service-line pages that already out-rank the nearest competitor in its strongest categories. The audit's job was to find where a strong footprint was quietly leaking its most valuable searches.

It was leaking in three places a dashboard would never surface:

1. **The money intent was the smallest slice.** Over half the footprint was informational — top-of-funnel content that earns impressions but books no one. The commercial and transactional terms that actually convert were barely a third of the keyword mix combined.
2. **Live pages were pointed at the wrong words.** At least one service page was optimized for a phrase with effectively zero monthly search volume. It looked fine in a report and brought in nobody. Meanwhile the practice's own blog and provider pages were out-ranking its money pages for the exact terms those money pages were built to win.
3. **The architecture was both the advantage and the gap.** The service-line pages were strong, but the map was incomplete — several of the ways customers actually search (by symptom, by procedure, by specialist, by “near me”) had no dedicated page to catch them.

None of this is a content-volume problem. It is an **intent, architecture, and internal-linking** problem — which means most of the fix is re-pointing and connecting assets the site already has, not writing from a blank page. The 90-day game plan at the end of this document sequences that work so the fastest wins land first.

30,000+

RANKING KEYWORDS
top-100 positions, at
baseline

100K+

ORGANIC CLICKS /
QUARTER
~1,200+ per day

~2×

SERVICE-PAGE
RANKINGS VS. RIVAL
in top categories

~1.5%

SITEWIDE CTR
below its position
benchmark

GSC → Performance (clicks, CTR), 3-month baseline • SEMrush → Domain Overview (ranking keywords)
• SEMrush → Organic Research (service-page position counts vs. nearest competitor). Figures rounded.

Current-state footprint

Where the practice stood the day we started — the baseline every later result is measured against.

Search performance (GSC, 3-month baseline)

METRIC	QUARTER	28-DAY	READ
Organic clicks	~110,000	~35,000	Strong volume; stable
Impressions	~7.0M	~2.5M	High visibility
Average CTR	~1.5%	~1.4%	Below benchmark (3%+ expected at pos 7–9)
Average position	~8.9	~7.0	Improving

GSC → Performance → Search results, URL-prefix property. Quarter = trailing 3 months; 28-day = trailing 4 weeks. Figures rounded.

The headline here is the CTR gap. The site earns roughly **2.5 million impressions every four weeks** but converts about 1.4% of them to clicks — roughly half of what its average position of ~7 should command. That gap is the single cheapest win in the account: the rankings are already there; the titles, metas, and conversion paths are leaving clicks on the table.

Keyword intent mix (SEMrush)

The clearest signal in the whole audit. Sort the ranking keywords by search intent and the problem is obvious — the biggest slice is the one that books no one.

INTENT	SHARE OF KEYWORDS	SHARE OF TRAFFIC	WHAT IT MEANS
Informational	57.7%	~50%	“What is...” / “how to...” — a feeder, not a booking
Commercial	~18%	~18%	Comparison / provider-shopping intent — grow this
Navigational	~12%	~24%	Brand searches — healthy
Transactional	~12%	~8%	Ready-to-book intent — smallest slice, biggest opportunity

SEMrush → Organic Research → keyword intent breakdown, audit-cycle pull. Percentages rounded; shares do not sum to 100 due to rounding and multi-intent terms.

THE TAKEAWAY

Informational content is doing its job — it is a huge top-of-funnel audience. But the terms that turn into appointments (commercial + transactional) are barely 30% of the footprint combined. The plan is not to abandon the informational content; it is to **wire it to the commercial and transactional pages it was always pointing at**, and to go win the money terms directly.

Authority & competitive standing (SEMrush)

DOMAIN REPUTATION

Authority Score: mid-40s. Ahead of some regional rivals, behind the category’s national leaders (mid-50s to ~70). Referring domains sit around 2,000 — a solid base, roughly 30–60% behind the strongest competitors.

WHERE IT WINS & LOSES

Service-line pages out-rank the nearest competitor ~2× in the top categories — a real, defensible strength. The competitor pulls ahead only on **breadth of supporting content**, out-ranking the practice by up to ~8× on one category’s long-tail cluster, driven almost entirely by blog posts the practice hasn’t written yet.

SEMrush → Domain Overview (Authority Score, referring domains) · SEMrush → Organic Research, position counts by URL folder vs. nearest competitor. Figures rounded to ranges. “Off-site authority” = domain-reputation growth, not a paid tactic.

Technical findings

The technical layer was, on balance, healthy — which is good news, because it means the growth work isn’t blocked behind a rebuild. But four specific issues were quietly taxing the whole site.

Indexation & crawl

SIGNAL	STATE AT BASELINE	VERDICT
Indexed pages	~1,000 of ~3,000 known URLs	Expected for a large CMS with heavy redirect/canonical normalization
Crawled — currently not indexed	Growing (~80 → ~110 over the cycle)	Flag. A rising bucket means Google is choosing not to index new/thin pages — investigate the cause before publishing more
Legacy 404s	~100 URLs, 0 clicks / 0 impressions in 90 days	“Let them 404.” No redirect map needed — verified zero equity to preserve
Pages with redirect / alt-canonical	High, but all www/protocol + old-CMS normalization	Expected. No action

GSC → Indexing → Pages report, audit-cycle snapshot. 404 equity confirmed via GSC Performance (per-URL clicks/impressions, trailing 90 days). Counts rounded.

FINDING — PARAMETER POLLUTION

Tracking parameters appended to internal links (from a third-party review/social tool) were creating a **duplicate, separately-indexed URL for every provider and location page** — the clean URL and the ?utm_... URL both in Google’s index, splitting link equity and inflating the “competing pages” count for the site’s own keywords. Fix: canonical tags on the parameterized URLs pointing to the clean version, and strip the parameters before they reach the server. This is invisible in any dashboard and material to rankings.

Core Web Vitals & speed

Pass

MOBILE CWV (FIELD)
all URLs “good”

Pass

DESKTOP CWV (FIELD)
all URLs “good”

0

POOR / NEEDS-WORK URLS
at baseline

Chrome UX Report field data via GSC → Core Web Vitals · Lighthouse lab audit on representative templates. Snapshot at audit cycle.

Speed is a genuine strength here — a clean bill of health on Core Web Vitals removes the most common technical drag on rankings. (On sites where lab LCP is weak even while field data passes, we flag it as a “CWV time bomb” and schedule a proactive fix; that was not the case here.)

Schema markup

SCHEMA	COVERAGE AT BASELINE	VERDICT
Organization / WebSite	Present sitewide	Good foundation
Provider / Person schema	Present on top provider pages	Strength — ahead of most category competitors
Service schema on service pages	Absent on the ~9 core service pages	Major gap. The nearest competitor marks up every service page — a direct rich-result and AI-answer advantage
Review / AggregateRating	Not surfaced on service or location pages	Missed. A strong review corpus that Google can’t read in structured data
BreadcrumbList	Not implemented sitewide	Gap — forfeits breadcrumb rich results and cedes architectural clarity to crawlers
LocalBusiness · MedicalWebPage	Unverified / absent at baseline	Verify LocalBusiness on location pages in Rich Results Test; add MedicalWebPage to the YMYL condition/ procedure pages

Live on-page DOM crawl of representative templates • Google Rich Results Test. All schema deployed through the CMS’s native fields / SEO plugin – never hand-injected code.

Internal linking & conversion paths

FINDING – A HALF-MILLION IMPRESSIONS WITH NO EXIT

The top 10 blog posts by impressions — carrying **500,000+ combined impressions** — had **no scheduling CTA and no internal link** to the relevant service page. The site was earning enormous top-of-funnel attention and then giving those readers nowhere to go. Fix: a mid-article CTA, a provider/ service card, and a contextual internal link on every high-traffic informational post — turning a content library into a funnel.

On-page & keyword architecture

This is the core of the audit — and the part no automated tool produces. The question isn’t “are the pages optimized?” It’s “is there a page built to catch every distinct way a customer searches, and is each one pointed at a term people actually use?”

The search-entrance model

Customers don’t search one way. For a specialty practice, the same person might arrive through any of five distinct doors — and each door needs its own dedicated page, targeted at that door’s intent and volume:

ENTRANCE	HOW THEY SEARCH	PAGE THAT SHOULD CATCH IT
Condition	the diagnosis by name	Condition page (informational → feeds the service page)
Symptom	what they feel, before they have a diagnosis	Symptom page (high-volume, top-of-funnel)
Procedure	the specific treatment or surgery	Procedure / service page (commercial + transactional)
Specialist	the type of provider	Specialty hub + provider pages (commercial)
“Near me”	service + location	Location page (local intent)

Model applied against SEMrush keyword universe + GSC query data for the practice’s categories. Each entrance verified to have real search volume before a page is recommended.

Audited against this model, the practice's strongest service line was nearly complete — but most others were missing their symptom and procedure doors entirely. Those gaps are the build list for Month 2, prioritized by verified volume and difficulty.

The catch: a page pointed at a phrase nobody searches

FINDING

One live service page was optimized end-to-end for a phrase with **effectively zero monthly search volume**. It read well and validated clean in every tool — it simply described the service in words no customer types. Pulling the page's own GSC query data surfaced the term real customers do use for that service (a few hundred searches a month, clear commercial intent). Re-pointing the title, H1, and body to the real term is a one-hour change that turns a dead page into a live one. **This is the clearest case of what a hands-on audit surfaces that an automated dashboard never will.**

Combine vs. separate — a data-backed rule, not a preference

The recurring architectural question is whether near-synonym topics get one page or several. We decide it with data, every time, not with a house style:

- ✓ **Keep separate** when each variant already holds real, measurable impression equity of its own and the live SERP shows dedicated pages winning. Merging would sacrifice equity you can see in GSC for a theory that a combined page ranks better.
- ✓ **Combine / redirect** when you find a cluster of thin, near-duplicate pages splitting equity among themselves — in this account, roughly **30 near-identical service URLs** with essentially the same body content were consolidated into their canonical hubs, concentrating link equity instead of scattering it.

Decision inputs per topic: GSC per-URL impressions (equity to protect) + live SERP composition (are dedicated or combined pages ranking?) + SEMrush volume/difficulty. Rule applied per cluster, evidence recorded.

The geo-quarantine rule

SERVICE / PROCEDURE PAGES

No city modifiers. These pages target the highest-volume national head terms for the procedure. Bolting a city name onto them caps their ceiling at local volume and competes with the pages that should own local.

HOMEPAGE & LOCATION PAGES

Geo lives here. The homepage and the per-location pages carry the geographic targeting — “[service] + [city]” and “near me” intent — so the practice competes locally without diluting its national procedure pages.

Rule verified against SERP intent per query: national procedure terms return national/service pages; “[service] [city]” terms return local/location pages. Keeping the two page types on separate keyword targets stops them cannibalizing each other.

Cannibalization: the site competing with itself

Because the money terms had no clear single owner, the practice’s own pages were fighting each other in the results — the classic symptom of the architecture gaps above.

MONEY QUERY	PAGES COMPETING	WHERE THE SERVICE PAGE RANKED	WHAT OUTRANKED IT
Category A procedure	~90	pos ~12	A provider page (pos ~3) & a blog post (pos ~5)
Category B procedure	~85	pos ~24	A patient-story blog (pos ~9)
Category C condition	~30	pos ~35	A tangential internal page (pos ~7)

GSC → Performance → filter by query → Pages tab, 28-day window. Each row = multiple URLs from this site ranking for one query, diluting position and CTR. Positions rounded.

The fix is rarely to delete the competing page — the blog posts often earn real traffic. It is to **designate one owner per term**, strengthen it, and point the others at it with internal links. Where a blog legitimately owns an informational query, we leave it and simply add the conversion path.

THE COMPETITOR GAP

Even, then not — the gap is content breadth

Head-to-head on the pages that matter, the practice is winning. The gap opens only in one place: the volume of supporting content feeding the competitor’s money pages.

~2×

SERVICE-PAGE RANKINGS

practice leads, top categories

~8×

COMPETITOR’S LONG-TAIL
EDGE

one category, blog-driven

~1,700

BLOG-DRIVEN POSITIONS

the practice hasn’t written yet

Read: the practice wins the commercial pages and loses the informational long tail that feeds them. That is a closable gap — a content-build list, not a structural disadvantage. And it’s the reason Month 3 turns toward depth and off-site authority once the money pages are fixed.

SEMrush → Organic Research, position counts by folder for the practice vs. its nearest competitor, one category taken as the worked case. Figures rounded.

Prioritized fixes

Everything the audit surfaced, ranked by impact against effort, so the sequence is obvious. High-impact / low-effort items run first — those are the “quick wins” that fund the trust to do the deeper work.

ISSUE	IMPACT	EFFORT	WHEN
Title/meta rewrites on already-ranking, under-clicked pages	High — can multiply CTR with no ranking change	Low	Month 1
Re-point the zero-volume page to its real commercial term	High — revives a dead page	Low	Month 1
Add scheduling CTA + internal links to top informational posts	High — funnel for 500K+ impressions	Low–Med	Month 1
Canonicalize parameter-polluted duplicate URLs	High — stops equity split sitewide	Med (dev)	Month 1
Consolidate ~30 near-duplicate service URLs	Med — concentrates link equity	Low	Month 1
Service + Review + BreadcrumbList schema across core pages	High — rich results + AI-answer eligibility	Med	Month 1–2
Resolve cannibalization — one owner per money term	High — unifies diluted rankings	Med	Month 2
Build the missing symptom & procedure pages (search-entrance gaps)	High — captures unserved demand	High	Month 2–3
Expand thin service pages to competitive depth	Med–High — depth signals	High	Month 2–3
Supporting-content build to close the long-tail gap	Med — feeds money pages over time	High	Month 3+
Off-site authority building (domain reputation)	Med — lifts the whole domain	Ongoing	Month 2+
Verify LocalBusiness schema on location pages	Low–Med — local eligibility	Low	Month 1

Impact = expected effect on clicks / bookings, weighted by the page's existing demand (GSC impressions). Effort = KDM execution time; "(dev)" marks items needing a developer request. Sequencing favors high-impact / low-effort first.

The 90-day game plan

The audit is only useful if it becomes a sequence. This is the plan the findings produced — not a stock checklist — ordered so the fastest wins land first and each month unblocks the next.

Month 1 — Foundation & quick wins

Verify, fix the cheap high-impact leaks, and stand up measurement. Most of this is re-pointing assets that already exist.

- ✓ Full technical crawl — verify indexation, schema coverage, speed, and internal-link structure against the live site (never against stale audit data).
- ✓ Title/meta rewrites on the highest-impression, lowest-CTR pages — each option presented with keyword volume and pixel-width before it ships.
- ✓ Re-point the zero-volume service page to the real commercial term pulled from its own query data.
- ✓ Add a scheduling CTA, a service card, and a contextual internal link to the top informational posts — a conversion path for 500K+ impressions.
- ✓ Consolidate the ~30 near-duplicate service URLs into their canonical hubs.
- ✓ Submit the developer request to canonicalize the parameter-polluted duplicate URLs.
- ✓ Tighten Google Business Profile + review posture for the map pack.
- ✓ Wire GA4 + Tag Manager so every call, form, and booking action is tracked from day one — you can't prove ROI on what you don't measure.
- ✓ Establish the baseline report every later month is measured against.

Month 2 — Build the map

With the leaks closed and measurement live, build the pages the search-entrance model says are missing — highest-value doors first.

- ✓ Deploy Service and Review schema across the ~9 core service pages, BreadcrumbList sitewide, and MedicalWebPage on the condition/procedure pages — through the CMS's native fields, for rich-result and AI-answer eligibility.
- ✓ Resolve cannibalization: designate one owner per money term, strengthen it, and point the competing pages at it with internal links.
- ✓ Build the missing symptom and procedure pages — each targeted at a term with verified volume and matching SERP intent.

- ✓ Expand the thinnest money pages toward competitive depth (targeting the depth the leading competitor demonstrates).
- ✓ Wire internal links cluster-to-hub as pages are authored, so authority compounds instead of pooling.
- ✓ Begin off-site authority building to grow domain reputation, pointed at the freshly strengthened money pages.

Month 3 — Compound

The foundation is fixed and the map is filling in. Now push depth, capture the clicks already earned, and start closing the long-tail gap.

- ✓ Content depth on the winning pages; begin the supporting-content build that closes the competitor's long-tail advantage.
- ✓ Lift CTR on pages ranking but under-clicked — the fastest, cheapest gains, now measurable against the Month 1 baseline.
- ✓ AI-answer optimization so the practice surfaces inside AI Overviews, ChatGPT, and Perplexity for its category.
- ✓ Continue off-site authority building at a steady cadence.
- ✓ Report in the conversion actions we can track — leads, calls, and bookings — not just keyword counts.

HOW WE'LL KNOW IT'S WORKING

Measured against the baseline, not against a promise

Every change with expected measurable impact gets a 30-day follow-up trigger against the Month 1 baseline. The metrics that matter, in order: booked calls and form submissions first; then CTR and average position on the money terms; then footprint growth. Market demand is tracked alongside — flat demand plus rising leads is the proof the engine is working, not the season.

CTR

PRIMARY EARLY SIGNAL
money terms, vs. baseline

Pos.

MONEY-TERM AVERAGE
cannibalization resolved

Leads

THE METRIC THAT PAYS
calls + form submissions

Targets set per-metric against the Month 1 GSC baseline • lead volume from the practice's call-tracking / CRM • market demand from Google Trends for the category + geography. Follow-ups created only for changes with expected measurable impact.

kava■

WHERE YOU'D START

This one's redacted. Yours would be about your business.

This is the depth of what your first month with Kava delivers — a full audit of your site and the 90-day plan that comes out of it, specific to your business and yours to keep. Want a look before that? On a quick intro call we'll walk you through a high-level version, live.

WHAT THIS IS

Your first month of work

BEFORE YOU COMMIT

A high-level look on an intro call

START

kavadigitalmarketing.com