CLIENT RESEARCH BRIEF · WORKING DRAFT

GEOLOCAL Research Brief

Date: 2026-10-01
Branch: planning-selfserve-audit
Opening: Three plain claims for first-time readers — no industry jargon — then evidence receipts.
Relationship: Append-only evidence layer for the alternate business proposal. Does not replace the product architecture source of truth.

CONFIDENTIAL · WORKING DRAFT

RESEARCH BRIEF

Title: GEOLOCAL Research Brief — evidence behind North Star Presence + Recommendability

Date: 2026-10-01 (three claims rewritten same day for first-time / non-industry readers)

Branch: `planning-selfserve-audit`

Relationship: Append-only evidence layer for this alternate proposal. Does not replace the concept SoT (`SELF_SERVE_AI_DISCOVERABILITY_CONCEPT.md`). Labels in the evidence sections below: **practitioner-reported**, **vendor**, **primary study**, **UNVERIFIED**.

Three claims (read this first)

1. What the problem is.

A small business owner updates their Google listing — new phone number, new weekend hours — and thinks customers will find the right info everywhere. They will not. The same details still sit wrong or missing on Yelp, Bing, Apple Maps, and dozens of other online directories. A handful of big data companies quietly copy that stale information around the web. When someone asks an AI assistant — ChatGPT, Gemini, or anything like them — “where’s a good bakery near me?”, that assistant often pulls from those scattered, outdated copies. Sometimes the shop never shows up. Sometimes it shows up with last year’s number. The owner never sees the miss. Customers lose trust. Fixing one listing is not the same as being findable and accurately described wherever people discover local businesses today, including when they ask AI.

2. What the industry looks like today.

Most of the “help your local business get found” industry still chases Google Maps wins: polish the Google business profile, collect reviews, post photos, clean up a few directory listings. That work can bring more phone calls from Google. Companies sell subscription tools that push your listing to many sites at once, score how “complete” your profile looks inside their own software, and track where you rank on a map. Almost none

of those products answer the question owners now care about: what do AI assistants actually say about my business — do they find me, recommend me, and get the facts right? Outside research shows AI assistants name far fewer businesses than the short list of three Google shows at the top of many local searches, disagree with each other, and often get hours or phone numbers wrong even when they recognize the shop. People who work in this field warn that big listing subscriptions can lock up your data — stop paying and your listings may disappear. Small owners who do the work themselves usually get the credit for results. Customers are asking AI for local recommendations more than they did a year ago, but the tools for sale still mostly optimize last year's game.

3. What we can build to solve it — and why it can be a business.

We build a product that starts with what the owner can feel: ask the major AI assistants what they say about this business, then show what is missing, wrong, or stale. That check is the reason to sign up. Behind it, one action files corrections on the places that actually feed those answers — the big data companies that fan out into many directories, and the major sites people and machines trust (Google, Apple, Bing, Yelp, and similar) — so the owner does not fill out a dozen forms. We still help them keep Google strong, because Google is where many customers double-check after an AI suggestion, and because a correct Google listing is useful proof when we ask other sites to update. We measure success simply: after the fixes, ask the same AI questions again — does the business show up, and is the description accurate? That is a viable business because owners already pay for listing tools that do not answer that question, AI recommendations are rising fast, and the gap between "I fixed Google" and "AI gets me right" is painful, visible, and mostly unserved.

The sections below keep the longer framing, a muffin-shop story, the research tables, and the detailed whys and hows. They are the receipts. The three claims above are the front door for anyone new to this industry.

North Star (how to read this brief)

The North Star for GEOLOCAL is **agent-era local SMB discoverability**: local businesses **discovered and recommended by AI assistants**. That intent (from `BUSINESS_PLAN.md`) is scored on this branch as two outcomes:

- **Presence** — being present and accurately represented wherever a local business is discovered (majors, umbrellas, directories, and the copies assistants lean on).
- **Recommendability** — would an AI pick you for a real buyer ask, and would it **get you right** (not just mention you)?

Everything below is evidence gathered **in service of that North Star** — what practitioners and vendors say moves Maps/calls vs what actually bears on AI Presence and Recommendability, what is waste, and where the proof is thin. It is not a general local-SEO encyclopedia; read every finding against "does this help an owner become discoverable and accurately recommended by agents?"

In plain English (the muffin shop)

Imagine Maya runs a small muffin shop. She updates her Google listing with the new phone number and weekend hours. She thinks she is done.

She is not done. Yelp still shows last year's number. Bing Places still shows last year's hours. A dozen smaller directories — and the big data companies that quietly feed them — still have the old copy. When someone asks an AI assistant "where's a good bakery near me?" that assistant often leans on those scattered copies. Sometimes Maya does not show up at all. Sometimes she shows up with the wrong number. Customers lose trust; Maya never sees why.

The pain, in human terms: fixing Google is not the same as being findable and accurately described wherever people — and AI — look.

The fix, in human terms: one button that (1) checks what AI assistants actually say about the muffin shop, (2) finds every stale copy, and (3) files the correction on each platform automatically. Maya never fills out a directory form. She confirms the truth once; the product does the chasing.

How we apply that to the product:

1. 1. **Lead with the AI-accuracy audit** — the hook Maya can feel. "Here's what ChatGPT / Gemini / others say about you — found, recommended, or wrong." That is Presence and Recommendability made visible.
2. 2. **Use the ticket set as the engine** — Lane A umbrellas (the data companies that fan out into many directories) plus Lane B majors (Google, Apple, Bing, Yelp, and the other big claim surfaces). The audit finds the stale copies; the tickets fix them with Maya's claimed truth (and Google as proof when it helps).
3. 3. **Measure success the way the North Star asks** — is the muffin shop findable, and is it accurately represented? Re-ask the same AI questions after corrections. Do not celebrate a directory "completeness" score that Maya's customers never see.

Keep Maya in mind while reading the evidence below. The tables and studies are the receipts behind this story — not a different product.

Pain

Local owners and agencies are stuck between two scoreboards that do not agree.

1. 1. **Maps / pack pain is familiar and call-dense.** Practitioners report large call and pack swings from GBP completeness, primary category, review velocity, and citation cleanup — but those wins are often Maps-lane attributions, not proof that ChatGPT / Perplexity / Gemini will name the shop correctly.
2. 2. **AI pain is selective and accuracy-fragile.** Primary studies (SOCi LVI 2026, Insites Jul 2026, BrightLocal Local AI Visibility Study Sep 2026) show assistants recommend far fewer locations than the 3-Pack, disagree with each other, persist only ~half the time on re-prompt, and frequently misstate NAP even when they "know" the business.
3. 3. **Owners cannot see the gap.** Localeze Completeness, citation-count tools, and vibe-check prompt grids do not measure Presence (are you on the claim surfaces assistants lean on?) or Recommendability (would an assistant pick

you *and* describe you accurately?). Selling “cleanup ROI” as if aggregator correction alone moves AI answers lacks published pack/call studies.

Practitioner snapshot (r/localseo; full thread fetch often blocked — treat as **practitioner-reported**, not replicated):

Claim (label)	Source thread (reddit)	Notes
~73% of shops invisible to 4+ of 6 AI engines (50-business observational audit)	1slake5	**Practitioner-reported** observational audit; small n; methodology not independently reviewed
15+ consistent listings claimed ~3× more likely to be recommended	same cluster / community commentary around 1slake5	**Practitioner-reported**; correlation ≠ causation
GBP completeness 82→100% → calls 14→26/mo (~86%), visibility 40→70% in ~12 days	1r23zji	**Practitioner-reported**; community pushback: may be reclassification / seasonality, not pure completeness lift
Primary category swap: page 2 → top 3 in ~3 weeks	1viqf8z	**Practitioner-reported** Maps/pack outcome
Delete ~40 bad citations: accuracy 74→98%, calls doubled ~6 weeks	1v7s2ae	**Practitioner-reported**; no controlled AI re-prompt evidence in thread
GBP opt ~70% calls ~2 months; AI mentions stayed flat (comment)	1o08abp	**Practitioner-reported**; useful as “Maps ≠ AI” anecdote
Reviews 11→34, page 2→pack, calls doubled ~6 weeks; attribution fight (~80% reviews?)	1vn7vjv	**Practitioner-reported**; heavy confounding
HVAC: review velocity as biggest lever	1vtlpki	**Practitioner-reported** vertical anecdote
Citations-only → top 8–15, not top 3	1m4woes	**Practitioner-reported**; citations help floor, not crown
Yext “data hostage” / listings vanish when stop paying; Synup cancel fee ~$1,980; Uberall enterprise-only	community threads (vendor pricing folklore)	**Practitioner-reported** / vendor-commercial; verify contracts case-by-case
DIY gets RESULT credit; paid pipes for scale past ~30 locations / 100+	community consensus	Product implication: ownership & tickets > hostage syndication
AI evidence thin vs Maps call evidence dense; ListingsAPI: no win case study	community meta	Aligns with BrightLocal / Insites gaps below

Pain for this proposal: the North Star fails when (a) the shop is missing from AI answers, (b) the shop is named but facts are wrong, or (c) the owner has no accuracy-aware audit + ticket path across Lane A umbrellas and Lane B majors.

What works (with evidence)

Maps / GBP / reviews (dense call evidence — mostly practitioner + expert surveys)

Lever	Evidence	Label	URL / cite
Primary GBP category, proximity, keywords in GBP title, address in city, open-now	Whitespark LSRF 2026 “pack kings” (survey posted ~Nov 6, 2025)	**Vendor / expert survey** (47 experts, 187 factors)	whitespark.ca/local-search-ranking-factors
GBP completeness / category / reviews → pack & calls	Multiple r/localseo case posts (table above)	**Practitioner-reported**	reddit threads linked above
Review volume & recency	Insites dual-surfaced avg **133.4** Google reviews vs **10.7** for neither; Whitespark: review recency underrated	**Primary study** + **expert survey**	Insites benchmark; Whitespark LSRF 2026
Dedicated service pages, best-of lists, industry domains, unstructured citation *quality*, review-site authority	Whitespark LSRF 2026 AI-factor discussion	**Vendor / expert survey**	same LSRF page
DIY ownership of GBP + majors	Community: DIY gets RESULT credit	**Practitioner-reported**	r/localseo meta

AI visibility (thinner causal evidence; stronger descriptive studies)

Finding	Evidence	Label	URL
AIO ~68% vs local packs ~39% on Whitespark local-business query set (540 queries; May 12, 2025)	Whitespark / Miriam Ellis prevalence study	**Primary study** (vendor-affiliated)	whitespark.ca case study
Maps mention 66% vs AI surfaces 32–38%; ChatGPT ∩ Google AIs ~19–20%; persistence ~50%; top citations GBP 28.6%, Yelp 9.5%, Facebook 2.2%	BrightLocal Local AI Visibility Study, Sep 16, 2026 (~200k prompts)	**Primary study** (vendor)	brightlocal.com/research/local-ai-visibility-study
Consumer AI local recs 6%→45% YoY; 23% used AI in last local search; of AI users ~18% contact direct, ~75% check Google/Maps	BrightLocal LCRS 2026 / related AI-trust cuts	**Primary study** (survey)	LCRS; LCRS AI trust
ChatGPT recommends **1.2%** of brand locations vs **35.9%** in 3-Pack (~30× selective); ~45% retail overlap traditional vs AI	SOCi LVI 2026 (~350k locations / 2,751 brands)	**Primary study** (vendor)	soci.ai news; LVI PDF
GBP detail match ChatGPT ~55.58% / Perplexity ~48.05%; never recommended unbranded 40.88% / 68.79%; “corroboration, not optimisation”; schema/llms.txt barely correlate	Insites 10k US locals (Jul 2026)	**Primary study** (vendor platform)	Insites benchmark; schema piece

What “works” for North Star Presence + Recommendability (synthesis, not a causal claim):

- **GBP excellence remains mandatory** — proof hub, Maps call volume, and the strongest public corroboration signal assistants echo.
- **Consistent, high-quality presence on majors + umbrellas** supports corroboration (Insites framing) more than vanity citation spam.

- **Review velocity and volume** correlate with both pack and dual-AI surfacing; still confound-heavy for AI alone.
- **Measuring AI separately from Maps** is itself a “what works” practice (BrightLocal Maps 66% vs AI 32–38%; SOCi selectivity).

What doesn't / waste

Anti-pattern	Why it fails / is waste	Label
Recurring listing fees for ≤~10 locations	Whitespark commentary: prefer one-time / DIY at small scale	**Vendor / practitioner**
Mass low-DA citation blasts	Expert consensus: noise; citations-only stalls at ~top 8–15	**Expert survey + practitioner**
Selling citation-cleanup-only as pack/AI ROI	**No clean published cleanup-only → +X% pack/calls study** (Whitespark: no clean citation-cleanup-only ROI study; BrightLocal: no published aggregator-correction → +X% pack/calls study)	Gap — do not sell
Yext-like ongoing pipe as the product	“Data hostage” / listings vanish when stop paying; Synup cancel-fee folklore (~$1,980); Uberall enterprise-only	**Practitioner-reported** commercial risk
Assuming Uberall breadth / older directory studies = cleanup ROI	Older breadth studies ≠ correction ROI	**Vendor historical** — not cleanup ROI
Treating Localeze Completeness Score as an AI audit	Completeness = hub fill-in (Primary / Discovery / Connect / Hours), not web AI Presence/Recommendability	Product fact (Lane A)
Schema / llms.txt as the AI lever	Insites: barely correlate with recommendation	**Primary study**
Citations-only strategy for top 3	Practitioner: top 8–15 ceiling without GBP/reviews/on-page	**Practitioner-reported**
ListingsAPI as proven win path	Community: no win case study found	**Practitioner-reported** / **UNVERIFIED** absence
Equating OpenAI↔Foursquare partnership with ranking lift	Partnership / Places adjacency ≠ proven ranking lever	Architecture caveat
Foursquare “60–70% of ChatGPT local”	Often phrased as “reports suggest”	**UNVERIFIED** — do not treat as fact

BrightLocal framing remains useful commercially: **Google first, LLMs second** — not because AI doesn’t matter (LCRS 6%→45%), but because Maps still mentions tracked businesses more often and most AI users still verify on Google/Maps.

AI visibility vs classic local SEO framing

Dimension	Classic local SEO (Maps / pack)	AI visibility (assistants / AIO)
Unit of success	Rank in 3-Pack / Maps; calls from GBP	Named in answer; facts correct; re-prompt persistence
Selectivity	3-Pack ~35.9% brand appearance (SOCi)	ChatGPT ~1.2% recommended (SOCi) — ~30× tighter
Overlap	Pack ranking helps but does not guarantee AI mention	BrightLocal: Maps 66% vs AI 32–38%; intersection ChatGPT∩Google AIs ~19–20%
Stability	Rankings move; calls are measurable	Persistence ~50% on re-prompt (BrightLocal)
Top “citations” / sources	GBP, site, reviews, citations	GBP still top among directories (~28.6%), but business sites dominate unique domains; Yelp/Facebook trail
Expert factor mix (Whitespark LSRF 2026)	Pack kings: primary category, proximity, GBP title keywords, in-city address, open-now	AI lean: best-of lists, service pages, industry domains, unstructured citation *quality*, review-site authority
Consumer behavior	Call / direction intent	LCRS: AI recs mainstream (45%); many still verify on Google/Maps
Accuracy failure mode	Wrong hours on GBP hurt conversion	Assistants misstate NAP ~half the time even when “aware” (Insites)
Product scoreboard	Rank + calls	**Presence + Recommendability** (this proposal’s North Star)

Framing rule for GEOLOCAL: classic local SEO is necessary infrastructure and the densest call evidence; AI visibility is a **separate, more selective** surface that needs its own audit (mentioned vs recommended vs **misstated**). Do not collapse them into one vanity “AI SEO” score.

Gaps and caveats

1. 1. **Causal AI ROI from cleanup is unpublished.** No BrightLocal / Whitespark / SOCi paper cleanly shows “aggregator correction alone → +X% AI mentions or pack.” Do not invent one.
2. 2. **r/localseo numbers are observational.** Threads like 1slake5 (73% invisible…), 1r23zji (completeness→calls), etc. are **practitioner-reported**; n is small; confounders (season, category change, ads) are real. Community already flags reclassification risk on completeness wins.
3. 3. **Vendor studies have incentives.** BrightLocal, Whitespark, SOCi, Insites sell adjacent products. Still useful as **primary study** descriptive baselines — cite methodology, don’t treat as gospel.
4. 4. **Foursquare 60–70% of ChatGPT local = UNVERIFIED** when phrased as “reports suggest.” OpenAI partnership ≠ ranking lever. Managed Users timing remains an open question in the proposal.
5. 5. **Localeze Lane A mechanics:** ~$99/year per location paywall; **no free suggest-edit**; Completeness Score = hub sections (Primary / Discovery / Connect / Hours), not a web AI audit; monthly partners often **30–90 days** — set owner expectations.
6. 6. **Uberall / Yext / Synup commercial caveats** are practitioner folklore until contract-checked; still inform “don’t build a hostage pipe.”

7. 7. **ListingsAPI:** no community win case study located — treat as non-proof, not proof of failure.
8. 8. **Insites "corroboration, not optimisation"** undercuts schema/llms.txt-first GEO pitches; aligns with ticketed accuracy across real claim surfaces.
9. 9. **SOCi sample is multi-location brands** (~350k locations). SMB single-shop dynamics may differ; still the clearest public selectivity ratio (~30×).
10. 10. **BrightLocal citation mix** (GBP 28.6% / Yelp 9.5% / Facebook 2.2%) is study-specific; website citations dominate unique domains — own site still matters.

Implications for this proposal (North Star Presence + Recommendability; Lane A/B)

Architecture already in this file / concept SoT:

- **Two-lane:** Lane A umbrellas (Data Axle, Localeze, Foursquare) + Lane B majors (GBP, Apple Business Connect, Bing Places, Yelp).
- **North Star:** Presence + Recommendability; `BUSINESS_PLAN.md` intent = discover + recommend by AI.

Build toward:

1. 1. **Accuracy-aware AI audit** — fixed buyer prompts; score Presence (found on claim surfaces / stomped copies) and Recommendability (named when it matters; facts match claimed truth). Prefer mentioned vs recommended vs **misstated** over vanity mention counts.
2. 2. **One-click tickets for Lane A and Lane B** — field / wrong / right / GBP-or-claimed-truth as proof. Umbrella-only leaves Core majors wrong; majors-only leaves Layer-1 feeds stomping the claim.
3. 3. **GBP excellence still required** — densest call evidence; strongest corroboration signal; proof hub for tickets. Do not market AI audit as a GBP substitute.
4. 4. **Don't sell fake cleanup ROI** — no published aggregator-correction → +X% pack/calls or AI-mention study. Sell measurable re-prompt improvement and accurate tickets, not guaranteed pack jumps.
5. 5. **Measure AI separately from Maps** — BrightLocal / SOCi / Insites all show divergence. Product dashboards must not blend pack rank into an "AI score."
6. 6. **DIY result credit ⇒ product posture** — make tickets and ownership easy; avoid Yext-like hostage pipes where Presence evaporates when the subscription stops. Paid pipes remain relevant for scale (~30+ / 100+ locations per practitioner consensus), as optional scale lanes — not the identity of the product.
7. 7. **Localeze / Foursquare honesty** — charge/time expectations (Localeze paywall, 30–90 day partners); Foursquare Placemaker suggest-edit is real; OpenAI adjacency is not a ranking claim; Managed Users later.
8. 8. **Evidence hygiene in GTM** — label practitioner vs primary study; flag UNVERIFIED (e.g. Foursquare 60–70%); keep this brief as the evidence appendix so the commercial ladder stays North-Star-tight.

Success check (unchanged intent): re-prompt the same asks; Presence and Recommendability improve (found, recommended, facts correct) — not "Localeze score went up," not "we submitted 400 citations."