

THE COMPLETE BLUEPRINT

The AI visibility blueprint for multi-location brands

A working system for getting hundreds of locations onto the AI search shortlist, and keeping them there. The governance model, the copy-paste schema, a 30-point audit you can run this week, and a phased rollout that does not depend on your dev team being free.

SOCIAL PLACES

FOR HEAD OFFICE & BRAND TEAMS

~35 MIN READ + WORKSHEETS

WHAT'S INSIDE

- 01 **Why AI visibility is a separate discipline** the shift, in one page
- 02 **The operating model** distributed brand governance, and who owns what
- 03 **The five layers** data, reviews, structure, content, governance
- 04 **The schema your web team needs** copy-paste Organization + LocalBusiness
- 05 **The 30-point location audit** score any location in 20 minutes
- 06 **The 90-day rollout** three phases, no dev bottleneck
- 07 **Measuring the gap** what to track, separate from Google

Why AI visibility is a separate discipline

Search stopped handing customers a ranked list. When someone asks an assistant where to eat, who to bank with, or which branch is open now, they get a shortlist of one or two names with a reason attached. The list is gone. The judging has been done for them.

This is not a refinement of local SEO. It is a different mechanism with different rules, and the brands that treat it as "SEO with extra steps" are the ones quietly falling off the shortlist. The numbers below are the case for treating it on its own terms.

~1%

of local business locations are recommended by the leading AI assistant, versus roughly 36% appearing in Google's local 3-pack.

6→45%

of consumers using AI to find local services in a single year, one of the fastest behavioural shifts on record.

45%

overlap between brands winning Google's map pack and brands appearing in AI answers. The rest is a blind spot.

The third number is the one that reframes everything. More than half of the locations winning Google today are invisible the moment a customer asks an AI instead. Whatever you have already built for traditional local search does not automatically carry over. AI visibility has to be earned separately, and it is estimated to be several times harder to achieve.

What an AI is actually doing

Google indexes pages and ranks them on relevance and authority. An assistant does something narrower. It assembles a confident answer from a small set of sources it trusts, then names only the businesses it can verify without risk. It is optimising not to be wrong, rather than to be complete. That single difference is why two things matter far more than they used to:

- **Reviews became a gate, not a leaderboard.** Below a certain rating or review-response rate, a location is filtered out entirely rather than ranked lower.
- **Anything unverifiable is dropped, not penalised.** Conflicting data across platforms, or detail buried behind scripts an assistant cannot read, does not score lower. It gets left out of the answer.

AI does not rank a weak location below a strong one. It declines to mention it. Invisibility is the default state, not the penalty.

02

The operating model: distributed brand governance

At one location, AI visibility is a checklist. At three hundred, it is an operations problem. Every requirement in this guide is public and well documented. The reason most multi-location brands still fail is not knowledge. It is that the work does not survive contact with scale.

A single business can hand-tune one set of structured data, one Google Business Profile, one review inbox. A brand with hundreds of locations cannot. The same task, repeated by different people on different timelines, produces drift. And drift is exactly what an AI reads as unreliability.

The structural answer is the same one that solved multi-location listings and reputation a decade ago, applied to a new front door. Head office owns the system, the data standard and the templates. Locations contribute genuine local detail within that structure. Neither extreme works. Fully centralised content is generic and reads as thin. Fully devolved execution drifts within a quarter.

Who owns what

The single most common failure is that nobody owns AI visibility, so it falls between marketing, web, and operations until it quietly degrades. Assign these four roles before anything else. They can be four people or, in a smaller team, one person wearing several hats, but the accountability has to be named.

HEAD OFFICE / BRAND LEAD

The system owner

Owns the data standard, the page templates, and the decision on what is centralised versus local. Accountable for the programme overall and for measuring the gap.

WEB / DEV TEAM

The structure owner

Implements server-side rendering and the schema templates once, so every new location inherits them automatically rather than being built by hand.

REPUTATION / CX TEAM

The gate keeper

Owns review generation and response cadence per location, and watches for any location drifting toward the exclusion threshold.

LOCAL OPERATORS

The detail contributors

Supply the genuinely local detail head office cannot: nearby landmarks, real hours, location-specific services and the questions their customers actually ask.

The test for any decision

When something changes at one location, ask: does the correction propagate automatically across every platform and page, or does it depend on someone remembering? If it depends on memory, it will drift. Everything in the next three sections exists to remove that dependency.

03

The five layers that hold at scale

Ordered deliberately. Each layer assumes the one before it is solid. Skipping ahead, schema before clean data, content before consistency, produces locations that look optimised and still get left off the list.

Layer 1 · Accurate, consistent location data everywhere

One source of truth for every location's name, address, phone, hours, categories and coordinates, pushed identically to Google, Apple, Bing and your own site. This is the floor. An assistant that finds conflicting facts about a location resolves the conflict by omitting

it. When structured data contradicts what is on the page, or what is on the Google Business Profile, the safest move for the machine is to trust none of it.

Doing this by hand across hundreds of locations is where most programmes drift. A listings management platform exists to push one source of truth to every platform at once and flag the moment a location falls out of sync.

Layer 2 · Reviews managed as a gate

AI treats sentiment as a threshold. Every location needs active review generation and, critically, response, because response rate and recency are themselves signals. The goal is not a higher average for the brand overall. It is no individual location falling below the line where assistants stop recommending it. One quiet branch with a slipping rating is a hole in the map, not a rounding error.

Tracking response rates and threshold drift across every location manually is the part that quietly slips. A reputation management platform surfaces which locations are approaching the line before they fall off it.

Layer 3 · Machine-readable structure on every page

Two parts. First, server-side rendering, so a location's core detail is visible without waiting for scripts to run. Assistants work inside a limited context window and often stop reading before client-side content loads. Second, the schema stack that AI parses. The full copy-paste version is in section 4.

If you do not want to build and maintain hundreds of structured location pages in-house, this is exactly what a managed store locator is for: server-rendered, schema-complete pages generated from a template per location. Social Places' AI Local Pages takes this further, auto-generating a hyper-local page per location and product, with LocalBusiness, Product and FAQPage schema built in.

Layer 4 · Genuinely local content, produced at scale

Distinct content per location: nearby landmarks, neighbourhood context, location-specific services and FAQs that answer the conversational questions customers actually ask an assistant. The hard part is not writing one good location page. It is producing hundreds that are each genuinely local rather than a city-name find-and-replace, which reads as thin and duplicative to both Google and AI.

Layer 5 · Governance, monitoring and measurement

The layer that turns the first four from a launch into a system. Who is accountable when a location drifts, how drift is detected before it costs visibility, and measurement of AI visibility itself, separate from Google rankings. Most brands have no tooling to even see this gap, which is covered in section 7.

04

The schema your web team needs

You do not need to write this yourself. You need to know it exists, know it matters, and hand it to whoever owns your website. This section is built to be forwarded.

The pattern for a multi-location brand is simple to describe and specific to implement: **one Organization block** declared once at the brand level, and **one LocalBusiness block per location page**, each linked back to the parent. The link is what builds entity trust, it tells an assistant that this branch in Sea Point belongs to a national brand it already recognises.

Block A: the brand, declared once

Lives on the homepage or About page. Declared a single time for the whole brand.

ORGANIZATION · SITEWIDE

JSON-LD

```
<script type="application/ld+json">
{
  "@context": "https://schema.org",
  "@type": "Organization",
  "@id": "https://brand.com/#organization",
  "name": "Your Brand",
  "url": "https://brand.com",
  "logo": "https://brand.com/logo.png",
  "sameAs": [
    "https://www.facebook.com/yourbrand",
    "https://www.linkedin.com/company/yourbrand",
    "https://en.wikipedia.org/wiki/Your_Brand"
  ]
}
```

```
}  
</script>
```

Block B: each location, one per page

Lives on every individual location page. The parentOrganization line and the per-location sameAs to that branch's own Google Business Profile are the two things most brands miss.

LOCALBUSINESS · PER LOCATION PAGE

JSON-LD

```
<script type="application/ld+json">  
{  
  "@context": "https://schema.org",  
  "@type": "Restaurant",           // use the most specific sub-type  
  "@id": "https://brand.com/locations/sea-point/#location",  
  "name": "Your Brand, Sea Point",  
  "parentOrganization": { "@id": "https://brand.com/#organization" },  
  "address": {  
    "@type": "PostalAddress",  
    "streetAddress": "123 Regent Road",  
    "addressLocality": "Sea Point, Cape Town",  
    "postalCode": "8005",  
    "addressCountry": "ZA"  
  },  
  "geo": {  
    "@type": "GeoCoordinates",  
    "latitude": "-33.9180",  
    "longitude": "18.3870"  
  },  
  "telephone": "+27 21 000 0000",  
  "openingHoursSpecification": [{  
    "@type": "OpeningHoursSpecification",  
    "dayOfWeek": ["Monday", "Tuesday", "Wednesday"],  
    "opens": "09:00", "closes": "21:00"  
  }],  
  "sameAs": [  
    "https://www.google.com/maps/place/?q=place_id:THIS_BRANCH"  
  ]  
}
```

HAND THIS TO YOUR WEB TEAM

Deploy Block A once, sitewide. Generate Block B from a template so every new location inherits it automatically, never hand-built page by page. Add Review and FAQPage schema to each location page on top of the above. Then validate every location with Google's Rich Results Test before launch. Match every field exactly to that location's Google Business Profile, because contradiction between the two causes the markup to be ignored.

The reference for the required and recommended fields is Google's own [LocalBusiness structured data documentation](#). If your locations are generated by a template the moment they are added, this is a one-time setup that then scales to any number of locations. If they are hand-built, it does not survive your next fifty openings, which is the entire point of the operating model in section 2.

The build-once-then-scale version of this is what Social Places [AI Local Pages](#) is built to handle as managed infrastructure, so the schema is generated and validated per location, and kept in sync as your data changes, rather than maintained by hand.

05

The 30-point location audit

Pick one location. Score it honestly against the 30 checks below, one point each. The result tells you whether that location is shortlist-ready, and the pattern across a sample of locations tells you where your programme is weakest.

Run it on your best-performing location and your worst. The gap between them is usually where the scale problem is hiding.

Location data / 8 points

- ☐ Name, address and phone are identical across Google, Apple, Bing and your own site.
- ☐ Trading hours match across every platform, including public holidays.
- ☐ The Google Business Profile category is the most specific available, not generic.
- ☐ Precise latitude and longitude are set, not just a street address.
- ☐ A correction made centrally propagates to every platform without manual re-entry.
- ☐ The location has no duplicate or unclaimed listings anywhere.

- ☐ The Google Business Profile is verified and actively managed.
 - ☐ Photos are current and specific to this location, not brand stock only.
-

Reviews and reputation / 7 points

- ☐ Average rating sits comfortably above the recommendation threshold (well above 4.0).
 - ☐ Reviews are recent, with a steady flow rather than a stale cluster.
 - ☐ Review response rate is high, not near zero.
 - ☐ Negative reviews are responded to, not just positive ones.
 - ☐ There is an active mechanism generating new reviews at this location.
 - ☐ Head office can see this location's response rate today, without asking the branch.
 - ☐ An alert would fire if this location drifted toward the exclusion threshold.
-

Page structure / 8 points

- ☐ The location has its own dedicated page, not a shared regional one.
 - ☐ Address, hours and services are visible without JavaScript running (server-side rendered).
 - ☐ LocalBusiness schema is present with the most specific sub-type.
 - ☐ GeoCoordinates are included in the schema.
 - ☐ The page links back to the parent brand via parentOrganization.
 - ☐ Review and FAQ schema are present on the page.
 - ☐ The schema passes Google's Rich Results Test with no critical errors.
 - ☐ Schema fields match the Google Business Profile exactly, no contradictions.
-

Local content / 7 points

- ☐ The page contains genuinely local detail, not a city-name swap of a template.
- ☐ Nearby landmarks or neighbourhood context are mentioned.
- ☐ Services or menu items specific to this location are listed.
- ☐ An FAQ answers the conversational questions a customer would ask an assistant.
- ☐ The content reads as written for humans, not stuffed for keywords.

- ☐ The page would launch fully populated on day one of a new opening, not months later.
- ☐ Content is refreshed on a cadence, not published once and abandoned.

Reading the score

SCORE	WHAT IT MEANS
26 to 30	Shortlist-ready. This location is structured the way assistants need. Hold the standard and watch for drift.
18 to 25	Partially visible. The fundamentals exist but gaps are costing recommendations. Prioritise the lowest-scoring group.
10 to 17	Mostly invisible to AI. Likely ranking on Google but absent from AI answers. Start with data and reviews before structure.
0 to 9	Invisible. The location is not structured for AI discovery at all. This is a programme, not a quick fix.

If your strong and weak locations score very differently, your problem is consistency at scale, not knowledge. That is the operating model in section 2, not a content task.

Running this across every location by hand is a real undertaking. If you would rather see it done for you, Social Places runs a listings and reputation audit across your locations, scoring AI search readiness alongside listings, reviews and local pages, and showing exactly where the gaps are.

06

The 90-day rollout

Sequenced so you are not blocked waiting for your dev team, and so each phase delivers visible progress before the next begins. The order matters: clean data first, because schema on top of inconsistent data just encodes the inconsistency.

1 Fix the floor

DAYS 1 TO 30 · OWNER: HEAD OFFICE + REPUTATION

- Establish one source of truth for all location data and reconcile it against every platform.
- Claim and de-duplicate any stray or unverified listings.
- Stand up review generation and a response cadence for every location, starting with the lowest-rated.
- Run the 30-point audit on a representative sample to baseline where you actually are.

No developer needed for most of this. It is data and process, and it is the work that makes everything after it count.

2 Build the structure once

DAYS 21 TO 60 · OWNER: WEB / DEV TEAM

- Move location pages to server-side rendering so detail is readable without scripts.
- Implement the Organization and LocalBusiness schema as a template, not page by page.
- Add Review and FAQ schema to the location-page template.
- Validate a sample of locations with the Rich Results Test and fix critical errors.

Overlaps with Phase 1 on purpose. The dev work can start while data cleanup continues. Built as a template, this scales to any number of locations.

3 Localise and govern

DAYS 45 TO 90 · OWNER: HEAD OFFICE + LOCAL OPERATORS

- Populate each location page with genuinely local content using a template that forces local detail.
- Set the governance cadence: who reviews drift, how often, and what triggers action.
- Stand up AI visibility measurement, separate from Google rankings (section 7).
- Re-run the 30-point audit and compare to the Phase 1 baseline.

By day 90 you have a system that runs, not a project that ended. Opening location three hundred and one becomes a structured step, not a fresh manual job.

Measuring the gap

A healthy Google report can hide a complete AI blind spot. Because only about 45% of map-pack winners also appear in AI answers, your existing local SEO dashboard is not telling you whether you are on the shortlist. You have to measure AI visibility on its own.

You do not need perfect data to start. You need a repeatable way to ask the questions your customers ask, and to track whether your locations come back. Track these, per location or per region:

WHAT TO TRACK	WHY IT MATTERS
AI recommendation rate	For a set of real customer-style queries, how often does each location get named by an assistant. This is the headline number Google cannot give you.
Profile accuracy in AI answers	When a location is named, is the detail the assistant states actually correct. Inaccurate is its own risk.
Review threshold proximity	Which locations are drifting toward the rating or response-rate line below which they get filtered out.
Silent locations	Which locations have gone quiet on reviews or updates, the early signal of a region falling out of answers.
Google vs AI overlap	The share of your Google-winning locations that are also appearing in AI answers. The gap is your exposure.

The question that justifies the whole programme

If a hundred of your locations silently dropped out of AI recommendations next month, how long until anyone at head office knew? If the honest answer is "we would not", measurement is where to start, before any of the fixes.

Turning this into a number you can watch is exactly what Social Places' AI Visibility is built for: one weekly score of how often assistants name each location, tracked by assistant, query and location, plus your share of voice against the competitors you name.

For the principles Google publishes on being surfaced in AI features, see Google's [guidance on AI features in Search](#).

WHERE THIS LEAVES YOU

Knowing is not the moat. Running it across every location is.

Everything in this guide is public. What separates the brands on the shortlist from the ones who are invisible is whether the work runs as a governed system across every location, or as a checklist someone hopes gets done. That is an operations problem, and it is solvable.

Social Places helps multi-location brands run exactly this system: one platform and one managed team keeping location data accurate, reviews managed, pages structured and visibility measured across hundreds or thousands of locations at once. You have the blueprint. The question is whether you run it yourself or have it run for you.

If you would rather see exactly where your locations stand before deciding anything, request a full audit. We score your AI search readiness alongside listings, reviews and local pages, so you start with the gaps, not a guess.

[Request a full audit →](#)

[Contact us for more →](#)



SOURCES & NOTES

Figures on AI recommendation rates, Google overlap and consumer adoption reflect 2026 third-party local search research, analysing several hundred thousand locations across thousands of multi-location brands. They are point-in-time estimates and are summarised here for context; consult the original research for full methodology before relying on any single figure. Schema patterns follow Google Search Central's published LocalBusiness and AI-features

documentation. The code samples are illustrative starting points, your web team should adapt fields to each location and validate before launch. AI search behaviour changes rapidly. Treat this guide as a working framework, not a fixed specification.

© 2026 Social Places. Provided for general information and does not constitute a performance guarantee.