# Local SEO Indexing: Checking NAP Citations and Location Pages

When you talk about local SEO indexing, checking NAP citations and location pages isn't something you can defer until after a site redesign or a merger. It’s the floorboards under your local presence — rotten planks and nobody stands. A single phone number drift across 15 directories can silently crater your Google Business Profile rankings because Google trusts what it sees third-party-wide, not just what you think you updated yesterday. The same applies to location pages on your own domain: if Google can’t find them, can’t parse the canonical, or decides the content is mostly boilerplate, the entire local pack pipeline goes dry.

Approximately 46% of Google searches carry local intent, and BrightLocal’s surveys repeatedly show that 68% of consumers would stop using a local business after encountering incorrect information online. That’s not a small leak. Yet I still audit brands spending $10k/month on ads while three of their store URLs return 404 or still list a forwarding number from three years ago.

This is not theory drawn from a dashboard. In practice, when you inherit a multi-location client whose expansion was driven by acquisitions, the NAP sprawl is often a historical mess: old Yext tiers, forgotten Yelp duplicates, franchise portals that autogenerate location pages nobody claimed. You don’t fix that with a dashboard widget. You go hunt down what’s actually indexed, and you compare the raw Google-visible text to a golden record.

## Why NAP Inconsistencies Hit Harder Than You Assume

Google’s local algorithm isn’t just a simple lookup. It cross-references signals from multiple sources — your own website, the GBP listing, structured data on the page, and a constellation of external directories. When the business name, address, or phone number differs even by a suite number or a formatted sting like “St.” vs “Street,” the entity graph weakens. Google’s systems start treating the variants as if they belong to multiple, less authoritative entities. That’s why a chain with 40 locations can see random “locations not appearing on Google Maps even though the page is indexed.” The indexing is there, but the confidence to show the Knowledge Panel side by side is not.

**Practical note:** Google's own documentation on structured data emphasizes that the LocalBusiness schema must match the real-world truth exactly. A mismatch between the name in your JSON-LD and the title on the page can cause the rich result to be suppressed without any crawl error.

I've seen a hospitality group where a single digit transposition in the phone number — an old VOIP-forwarding line — propagated into the schema of 27 location pages. Pages indexed fine. Local pack visibility? Gone for 11 stores within four days. The only clue was a drop in direction requests, not a Search Console alert. That's the kind of silent failure you're looking for.

## The Anatomy of an Indexed, Ranking-Ready Location Page

Indexing alone isn't enough. A location page must satisfy a lean set of "can it assemble the local entity" signals. If you treat a URL that returns 200 and shows up in site: as a win, you'll miss the deeper ailments.

Break it down. The page needs:
— A unique, crawlable URL with self-canonical and no accidental noindex in the HTTP header.
— Exact NAP text present in the visible HTML (not buried in an iframe or generated client-side only after user interaction).
— Matching LocalBusiness schema, referencing the same @id across the site and ideally linked to the GBP via sameAs.
— Store-specific operational content (hours, services, accessible entrance notes) distinct from the 20 other location templates.
— An HTML sitemap and XML sitemap entry that Googlebot can reach without hitting a chain of 301s.

Too many location pages are thin duplicates wearing a different city name in the H1. Google's quality guidelines — yes, the ones that mention doorway pages — absolutely apply here. If the only text change between /locations/dallas and /locations/houston is the city string, you're on thin ice, and the indexing might be partial or flaky: sometimes in, sometimes out.

## How to Audit NAP Citations Without Losing Your Mind

A real audit isn’t a spreadsheet with three columns. You need a repeatable pipeline that touches both the directory side and the owned location pages. I’ll show the compact version I use when time-boxed to two hours per brand.

```
flowchart LR
  A[Pull GBP listing + golden NAP] --> B[Export 20-50 known citation U
RLs]
  B --> C{Indexed in Google?}
  C -- Yes --> D[Scrape visible NAP text from cached snippet or page]
  C -- No --> E[Flag as unindexed; trigger re-
crawl or submit to Indexing API]
  D --> F[Compare NAP tuple to golden record]
  F -- Mismatch --> G[Log discrepancy + priority]
  F -- Match --> H[Mark verified]
  G --> I[Queue for correction & re-indexing signal]
```

Step one: build a golden NAP record from the GBP dashboard and the official corporate location database. That’s your source of truth, not the person who “thinks the phone number is up to date.” Step two: generate a target URL list. Include the obvious — Yelp, Foursquare, Yellow Pages, Facebook — but also pull citations from industry-specific directories (healthcare, legal, automotive) using a tool like Moz Local or white-label reseller APIs. Step three: don’t assume every URL is indexed. Use a bulk index check.

Here’s a curl snippet that pings Google’s URL Inspection API for a single page, assuming you have OAuth set up with the Indexing-capable scope:

```
curl -X POST \
  https://searchconsole.googleapis.com/v1/urlInspection/index:inspect
\
  -H "Authorization: Bearer $(gcloud auth print-access-token)" \
  -H "Content-Type: application/json" \
  -d '{
    "inspectionUrl": "https://example.com/locations/austin",
    "siteUrl": "sc-domain:example.com"
  }'
```

That endpoint returns the index state plus any AMP or mobile-usability issues. It

doesn’t tell you if the NAP is right — just that Google fetched it. For bulk checking hundreds of location pages, the Google API quota limits you fast; that’s when an external index-checker API with higher throughput becomes the actual working choice. SpeedyIndex’s bulk index checker processes thousands of URLs in one call — pragmatically, that saves me crossing the rate-limit cliff. Their API docs are straightforward enough to wire into a Python script in ten minutes.



Once you know what’s indexed, you still need to verify NAP correctness on the actual pages. I typically use a lightweight Python script that fetches the URL, extracts the visible text (via BeautifulSoup), and regex-matches phone number patterns and address formats, then compares against the golden record. It’s not perfect — dynamic JavaScript rendering is a different beast — but it catches 85% of mismatches. The edge case is pages where the phone number is an image or a click-to-call link with a obfuscated href. For those, I’ll manually spot-check the rendered cache.

```
import requests
from bs4 import BeautifulSoup
import re
golden_phone = "512-555-0199"
urls = ["https://example.com/locations/round-
rock", "https://example.com/locations/cedar-park"]
def check_nap(url):
    r = requests.get(url, headers={"User-Agent": "Mozilla/5.0"})
    soup = BeautifulSoup(r.text, "html.parser")
    text = soup.get_text()
    # simple regex for phone; false positives possible but rare
    found_phones = re.findall(r'\(?\d{3}\)?[\s.-]?\d{3}[\s.-]?\d{4}',
text)
    normalized = [re.sub(r'\D', '', p) for p in found_phones]
    return golden_phone.replace("-", "") in normalized
for url in urls:
    print(f"{url}: {'OK' if check_nap(url) else 'MISMATCH'}")
```

**A real gotcha:** if the page serves a 200 but the body is mostly a spinner waiting

for JavaScript to hydrate, the above script sees an empty div and reports “no phone”—while the human-readable view shows a perfect NAP. Always combine automated extraction with a headful browser (Playwright/Puppeteer) for dynamic pages, or fall back to manual spot-checks for critical locations.

## When Google Indexes the Wrong Location Page

This is far more common than mismatched NAP on directories. Your canonical signals point to /locations/austin-main, but Google indexed /locations/austin?ref=homeservices&sessionid=abc — a parameter-ridden URL generated by a local CRM integration. Or worse, it indexed the PDF version of a flyer that lists a different phone number. The symptoms: the page title and snippet in search show garbage, or the location appears in organic results but never in the local pack. That’s a strong indicator your canonical and hreflang (if you have multilingual locations) are being ignored because the page signals got fragmented.

Use the URL Inspection tool inside Search Console for the “wrong” URL and look at the “User-declared canonical” versus “Google-selected canonical.” If they differ, you have a canonicalization conflict. In multi-location setups, I’ve seen this happen when the CMS appends a trailing slash inconsistently, or when the location page template sets a default canonical tag that points to the parent /locations/ index instead of the specific URL. Fix the tag, then request indexing of the correct URL. For batch, a sitemap with <lastmod> and a ping to the Indexing API nudges Google to re-evaluate. SpeedyIndex, for large sets, can submit a clean list of corrected URLs quickly, making the re-crawl signal less painful to send.

> Rule of thumb: never assume a 200 status means the page passes local entity assembly. Always cross-reference the indexed URL set with your intended canonical list before you blame the directories.

## Real-World Audit: A Multi-Location Retailer

Let me walk through a concrete case. A regional plumbing company with 14 branches had a migration from a subdomain to a subdirectory architecture six months prior. They came to us because Google was showing only 7 of the 14 locations in the local pack, despite all pages being indexed per a simple site: search. The deeper problem: only 9 location pages were indexed under the

correct canonical path; the other 5 were indexed only as old subdomain URLs (still serving content via a wildcard redirect that Google hadn't fully processed). Worse, on three of those old indexed subdomain pages, the phone number displayed was the previous owner's — the redirect didn't update visible contact info.

We pulled the GBP data, built a golden NAP, ran a batch index check via SpeedyIndex on all 14 intended subdirectory URLs plus the 5 subdomain artefacts. Result: 5 canonical pages were indeed not indexed, and the 3 indexable artefacts carried wrong NAP. We forced removal of the old URLs via temporary 410s, fixed the canonical chain to point strictly to the subdirectory versions, updated the schema on each real page, and submitted all 14 to re-Index. Within five days, 13 out of 14 re-appeared; the 14th needed manual fetching because a template bug set a meta robots noindex only on the "Emergency Services" page. Yes, one rogue noindex made an entire location invisible for months.

## Local Indexing Questions That Keep Coming Up

- **Can I rely on Google Search Console's Index Coverage report for location pages?**
  Yes, but it won't tell you if the NAP is correct or if the page triggers local pack eligibility. It only says "submitted and indexed." Use it to spot crawling anomalies, then overlay your own NAP check.
- **How often should I re-audit NAP citations?**
  For active directories that accept user edits (Yelp, Facebook), every 60 days. For static aggregators, quarterly. If you've had a phone number change, audit within 48 hours and prioritize high-credibility sources first.
- **Does a noindex on a PDF location flyer hurt the associated HTML page?**
  Only if Google confuses the PDF as the primary canonical. Prevent this by ensuring your HTML page uses self-canonical and the PDF is either blocked in robots.txt or served with X-Robots-Tag: noindex header. I've seen a chain where Google chose the PDF as canonical because it was older and had more backlinks — a mess.
- **What if my NAP is identical across all directories but the location page still doesn't rank?**
  Check the proximity of the page to the user's search centroid — local rankings are heavily proximity-weighted. Also verify the page isn't a template duplicate. Adding a "From this location" map snippet, local photos, and genuine customer reviews on the page can nudge the needle.

**5-point NAP consistency checklist:**

- Phone uses the same formatting everywhere, including country code when relevant.
- Street address matches the exact USPS-verified format, including suite and ZIP+4.
- Business name lacks extra keywords or location descriptors not in the GBP.
- Every location page has a unique, descriptive title tag and an H1 with natural city inclusion.
- Schema markup references the same @id chain and includes a sameAs link to the GBP.

## Don’t Let Invisible Location Pages Cost You Revenue

There’s a sunk-cost fallacy in local SEO: teams will invest heavily in content and PR, then assume the location data layer “takes care of itself.” It doesn’t. Indexing without consistency just amplifies wrong information. A large-scale study by a local citation platform found that businesses with fully consistent NAP across the top 50 directories enjoyed roughly 30% more Google Maps views than those with even two mismatched entries. That’s revenue left on the table because of a typo nobody checked.

Your next move is simple, but it requires commitment: pick a golden NAP, audit your owned pages first, then the directory ecosystem. Use brute-force bulk index checks to confirm what’s live in Google’s view, and script the NAP comparison where possible. When you find a problem, fix it at the source (not just the GBP dashboard), and trigger a re-crawl — either via Google’s Indexing API or a high-throughput service. You’ll stop bleeding visibility within a week or two, and you’ll never again stare at a location that’s *indexed* but invisible. The local algorithm respects entity clarity; give it that.

---

### Sources


1. Google Search Central. "SEO Starter Guide." developers.google.com
2. Ahrefs. "SEO Basics." ahrefs.com

3. Google Search Central. "Search Essentials." developers.google.com
4. Schema.org. "Getting Started with Schema." schema.org
5. Search Engine Journal. "SEO Guide." searchenginejournal.com