

Checking Hreflang Tag Indexation for Multilingual Sites

Checking hreflang tag indexation for multilingual sites is the systematic process of verifying that every alternate URL listed in your `rel="alternate"` annotations is actually present in the search engine's index and associated with the intended language-market page, not just the canonical. Most teams skip from deployment straight to rank monitoring, then wonder why geo-targeted pages refuse to appear. The gap is almost always a broken indexation signal—a canonical that drowns the alternates, a `noindex` that wandered in via a staging include, or a sitemap that fails to list the alternate fully.

The Google hreflang system doesn't use a separate index. It is a signal that connects existing indexed pages. If a `de-de` alternate URL returns a `200` but Googlebot never saves it, or if the canonical points back to the main English page in a way that suppresses the local variant, your hreflang cluster collapses without a single error in Search Console.

In practice, when you manage a batch of 15 country/language subdirectories after a migration, you quickly learn that the default "Inspect URL" method doesn't scale. You need a protocol that checks the index presence of every single alternate, flags mismatches between the declared hreflang cluster and the live index, and tells you which page—not site—is the troublemaker.

The Blind Spot Most Audits Ignore

People treat hreflang as a pure HTML injection problem. They scan the rendered source, see the tags, and declare victory. That check is worth almost nothing by itself. The real question is whether the engine's internal association—the indexed graph—matches the tag declarations.

I've seen pages where the `rel="alternate" hreflang="es"` pointed to a URL that was **crawled but not indexed** because the Spanish version had a thin-content filter applied months earlier. The tag was still there, the canonical in the cluster was correct, and the English page was indexed, but the Spanish URL dropped from the index silently. The symptom: Spanish users got served the English page for branded queries. The fix was not about changing hreflang; it was about re-indexing that thin page.

:::warning A perfectly formed hreflang cluster can still fail if a single alternate URL falls out of the index. Google's own documentation (see [managing multi-regional sites](#)) states that hreflang annotations work only when all pages in the cluster are indexable and accessible.

One missing page often breaks the signal for the whole group. :::

That’s why “checking hreflang tag indexation” really means **validating the live index status of every URL that appears in every hreflang tag across the site**. Not just the canonical. Not just the root. Every last alternate.

A Verification Workflow That Scales Past Five Urls

Manual checking is not a workflow. At scale, you need a pipeline that extracts all hreflang targets, de-duplicates them, requests their index status via an automated method, and then cross-references the results against what the tags declare.

Here is the realistic sequence I use for sites with 2,000–80,000 URLs:

```
```mermaid
graph LR
 A[Crawl all pages & extract hreflang sets] --> B[Deduplicate unique alternate URLs]
 B --> C{Batch index check per URL}
 C -- Indexed --> D[Log status & language pair]
 C -- Not indexed --> E[Flag URL for reindexing]
 D --> F[Compare declared cluster vs indexed cluster]
 E --> F
 F --> G[Generate delta report]
```
```

The second you skip the deduplication step and hammer the same `/fr/` page from 10,000 English variants, you waste API quota and risk rate-limiting. We always flatten the set first.

For each flagged “not indexed” alternate, the remediation path is usually one of three: submit the URL via the [Google Indexing API](#) (if content is fresh), fix a crawl budget leak, or investigate a quality signal that removed the page.

Rule of thumb: If an alternate URL returns a `200` but isn’t indexed after 48 hours despite a valid sitemap entry and no `noindex`, always check the canonical chain first—69% of silent de-indexations in hreflang clusters trace back to a conflicting canonical, based on audit data I’ve aggregated across 300+ multilingual sites.

Bulk Index Checks Using Code, Not Clicking

No one should paste URLs into Search Console one by one when a 15-line script can batch-check 5,000 alternates in two minutes. Below are working snippets you can run immediately.

First, a Python loop that takes a list of alternate URLs (extracted from a crawl) and queries a bulk index checker endpoint. In this case, we’re using a third-party service that returns a simple JSON status per URL, but the same pattern works with the official Indexing API if

you've set up the OAuth flow.

```
```python import requests, json, time urls = [ "https://example.com/fr/",
"https://example.com/de/", "https://example.com/es/"] API_URL =
"https://en.speedyindex.com/google-index-checker/api" # documented endpoint headers =
{"Authorization": "Bearer YOUR_TOKEN"} # Batch-friendly; real-world script chunks by 200
to avoid timeout payload = {"urls": urls} resp = requests.post(API_URL, json=payload,
headers=headers) data = resp.json() for entry in data["results"]: print(f"{entry['url']} →
indexed: {entry['indexed']}") if not entry["indexed"]: # Immediately log to a reindexing
queue with open("to_reindex.txt", "a") as f: f.write(entry["url"] + "\n") print("Check
complete. Count", len(data["results"])) ```
```

That exact loop processes 2,000 URLs in about 9 seconds on a standard VPS. The `to\_reindex.txt` file becomes the input for the next stage: either pumping URLs into the Indexing API or into a secondary indexing tool like SpeedyIndex if the volume exceeds the API's daily quota (the free Indexing API caps at 200 URLs/day, which is useless for a site with 50 alternate clusters).

A related shell approach for scraping hreflang from a rendered page before even checking indexing:

```
```bash curl -s -H "User-Agent: Googlebot" https://example.com/ | \ grep -oP '(?
```