The Hreflang Magic Trick – It Routes, It Doesn’t Rank via @sejournal, @billhunt



In 2018, Glenn Gabe published one of the more puzzling SEO case studies of the decade. A client was canonicalizing UK and Indian hreflang URLs to a U.S. version, technically wrong, and against Google’s own guidance. Yet, the “wrong” URLs kept surfacing correctly in country-specific search results, even though Search Console showed them as not indexed. When Gabe asked John Mueller about it during a webmaster hangout, Mueller confirmed the behavior was real: Google could swap in the correct country/language URL at the SERP level, even when only one page in the hreflang cluster was actually indexed. It was treated as a curiosity, a “magic trick,” as Gabe called it, because Google never fully explained the mechanism behind it.

Eight years later, it got a name. In August 2026, Google’s Gary Illyes responded to a LinkedIn question about this scenario and described these URLs as alternate names. His

explanation gives us a useful mental model for something international SEOs have been reverse-engineering for years: discovery, canonicalization, indexing, scoring, and serving are separate parts of the process, and a URL doesn't necessarily have to pass through all of them independently to appear in search.

That distinction helps explain where hreflang actually fits in the process. More importantly, as search moves toward semantic retrieval and AI-generated answers, it raises a larger question for international websites: if multiple market pages contain essentially the same information, what makes any one of them uniquely valuable?

First, Google Has To Discover The Alternatives

Before there's a cluster to canonicalize, a canonical to score, or an alternate URL to serve, Google has to discover the URLs. This is why XML sitemaps, internal links, and hreflang annotations become particularly important during migrations, replatforms, and international rollouts. Google needs to discover the U.S., UK, Australian, Italian, and other versions before it can understand how those resources relate to one another. Once discovered and crawled, those URLs don't automatically become independently indexed documents. Google still has to determine whether they represent genuinely different resources or duplicate and near-duplicate versions of essentially the same content. Some may warrant independent indexing and scoring, while others may be clustered with another URL.

This is an important distinction because we often treat "discovered," "crawled," and "indexed" as stages a URL simply progresses through. The alternate-names explanation demonstrates that Google can retain knowledge about a URL and its relationship to other URLs even when that URL isn't

independently represented in the index.

Hreflang Is A Cross-Reference, Not A Ranking Signal

One of the most persistent misconceptions about hreflang is that it somehow creates ranking value. I have always found it more useful to think of hreflang as a cross-reference table that tells Google that this URL is the Italian equivalent of another URL; this one is intended for English speakers in Australia, and another is intended for English speakers in the UK.

Google can use those relationships when determining which URL to present to a particular searcher, but that doesn't mean hreflang created the ranking in the first place. Across tens of thousands of hreflang implementations, I've never seen a case where adding hreflang suddenly caused a page to rank higher because of the annotation itself.

If a U.S. English page is ranking in the UK before hreflang is implemented, adding the correct UK alternate doesn't magically give the UK page a new ranking score. What can change is which URL Google presents for the ranking it has already determined.

This is what I find particularly useful about Illyes' explanation of "alternate names." Indexing and scoring can happen on one member of a duplicate cluster while the serving layer selects an appropriate alternate URL for the user. Hreflang helps Google understand that relationship and select the appropriate representation. In that sense, hreflang routes rather than ranks.

The 38-Market Problem

A global shoe brand I've worked with illustrates the problem at scale. It had roughly 98 market websites, about 38 of which

were in English and nearly identical to one another. 22 even shared the same dollar-sign currency symbol. That meant dozens of URLs effectively carried the same product information, the same language, and often the same pricing presentation, differentiated primarily by domain, directory, or a handful of localized details.

From the organization's perspective, these were 38 different market websites. From a machine's perspective, however, many of them represented essentially the same resource. There is little reason for Google to independently store, process, and score 38 copies if they contribute almost no additional information.

Instead, it can cluster near-duplicates, choose a canonical representation, and retain enough information about the alternatives to present the appropriate URL when the user's country or language calls for it. This helps explain the "magic trick" Gabe documented. The alternate URL didn't necessarily need its own independent ranking score to appear because Google already understood its relationship to the URL that had been indexed and scored.

The serving layer could handle the rest.

Canonical Tokenization And The Cost Of Duplication

This becomes easier to understand if we think about canonicalization as an efficiency problem. Search engines operate at enormous scale, and there is little benefit in repeatedly storing and processing substantially identical information when those resources can be confidently clustered together.

One member of a duplicate or near-duplicate cluster can become the primary representation while Google retains relationships

to the alternatives. In traditional search, those relationships can then be used at serving time to present the most appropriate URL.

AI makes this more interesting because an AI system isn't limited to returning one document as a blue link. It can retrieve information, extract evidence from different sources, and assemble an answer. As a result, the unit we need to focus on shifts from the URL itself to the information it contains.

Instead of only asking which URL should rank, we increasingly need to consider which information is useful for answering the user's question, where the best evidence for that information can be found, and whether that evidence is actually relevant to the user's market.

Motoko Hunt explored this distinction recently, to say hreflang can help identify the intended language or market relationship, but that doesn't mean an AI system will consider the content itself the best evidence for a user in that market.

That distinction becomes especially important for international websites with large numbers of highly similar market pages. If the underlying information is substantially the same, the presence of separate URLs and correct hreflang annotations doesn't necessarily give an AI system a reason to retrieve one market version over another. That reason increasingly has to come from the content itself.

AI Systems Are Getting Better At Recognizing Equivalence

Microsoft provided an interesting look at this problem in its description of Web IQ, the retrieval and grounding infrastructure supporting its AI experiences. Its multilingual models are designed to represent semantically equivalent

information as close together as possible, even when that information appears in different languages.

That matters because hreflang represents an explicitly declared relationship. We tell the search engine that two resources are equivalents intended for different audiences. Semantic retrieval systems are increasingly capable of identifying that equivalence from the content itself. Two passages can be recognized as saying essentially the same thing even if the web admin never explicitly declares the relationship between them.

Microsoft also describes retrieval as focused on passage-level evidence rather than treating the entire URL as the only meaningful unit. That is an important change in how we should think about international content. If an AI system can recognize equivalent information across pages and languages and retrieve the specific evidence it needs, creating additional URLs containing substantially identical information doesn’t necessarily create additional useful knowledge. It may simply create more copies of knowledge the system already has.

This is where what initially looks like a technical hreflang discussion becomes a much larger international content issue.

Hreflang Can’t Make Content Local

Hreflang can tell Google that one page is intended for the UK and another for Australia, but it cannot make either page more relevant to those markets. If 20 market pages contain essentially the same product descriptions, claims, specifications, FAQs, and supporting information, adding perfect hreflang annotations doesn’t suddenly make those pages meaningfully different.

The organization may have created 20 URLs, but from a retrieval perspective it may have created one piece of knowledge with 20 different addresses.

This is why I think localization needs to move beyond the traditional definition of translating content and swapping currencies, phone numbers, spelling, or a few market references. The more useful question is what someone in a particular market needs to know that someone in another market does not.

Products may differ by market. Regulations, certifications, availability, warranties, delivery options, and pricing models may be different. Customers may use different terminology, have different use cases, evaluate different competitors, require different integrations, or make purchase decisions based on completely different criteria. Local teams may also have evidence, customer examples, regulatory expertise, and market knowledge that never makes it onto the website.

Those differences give a market page a reason to exist and give search and AI systems something unique to retrieve.

Give Every Market Page A Reason To Exist

None of this makes hreflang less important. If anything, Google's explanation of alternate names makes its role clearer. Hreflang establishes the relationship between alternate language and market URLs and helps Google understand which version is intended for which audience. What it cannot do is create the market differentiation itself.

The content has to do that.

This becomes increasingly important as retrieval systems get better at understanding semantic equivalence across URLs and languages. If the U.S., UK, Australian, Singaporean, and South African pages all say essentially the same thing, machines are becoming increasingly capable of recognizing that regardless of how many URLs, folders, domains, or hreflang annotations

surround the content.

The objective shouldn’t be to make those systems believe pages are different when they aren’t. International teams should instead look for the places where markets really differ and ensure those differences are reflected in the content.

That requires more than an SEO or localization team working independently. International SEO, localization, content, product, legal, regional marketing, and local market teams may all hold different pieces of the information needed to make a market experience genuinely useful.

The technical implementation establishes the relationship; the local knowledge establishes the difference.

That is ultimately what Google’s “alternate names” explanation helps clarify. The mystery Gabe documented wasn’t a bug or loophole. It exposed the differences among indexing, scoring, and serving long before Google gave us the terminology for them.

As search moves further toward semantic retrieval and AI-generated answers, the bigger lesson isn’t simply how Google can serve an alternate URL that wasn’t independently indexed. It’s why we should give that alternate page enough unique local value that both machines and customers have a reason to care about it in the first place.

More Resources:

Featured Image: Roman Samborskyi/Shutterstock

https://www.searchenginejournal.com/the-hreflang-magic-trick-it-routes-it-doesnt-rank/591979/