Answerova

Google GEO Research

Google GEO Technical Foundations White Paper

An implementation and release guide for engineering, SEO, and operations teams

November 2026

Read online or switch languages

Contents

An implementation and release guide for engineering, SEO, and operations teams

03 Eligibility for Google AI Search

04 Crawling, HTTP status, and indexing

05 Snippet controls and Search generative AI settings

06 Rendering and accessible buying evidence

07 URL architecture and multilingual coverage

08 Structured data and visible facts

09 Release acceptance and regression checks

10 Impressions, citations, and sales outcomes

11 Sources

Eligibility for Google AI Search

AI Overviews and AI Mode follow Google Search fundamentals. A page must be indexed and eligible to appear in Search with a snippet before it can be shown as a supporting link.

Audit discovery, crawling, rendering, indexing, snippet eligibility, and actual exposure as distinct stages. Google can discover a URL and still choose another canonical or defer indexing. Eligibility does not guarantee indexing, visibility, or citation. Record technical acceptance separately from Search performance.

Prioritize URLs that help a buyer make a decision, such as model specifications, process comparisons, quotation terms, and compliance evidence. Google's generative features can search related subtopics for a single question. Build around a complete buying task and link to the evidence that supports it. [1] [2] [3]

What to verify

- Record the intended URL, Google-selected canonical, index status, and preview policy.
- Google requires no llms.txt file, AI-specific Schema, or set passage length. Keep those out of the Google release criteria.

Crawling, HTTP status, and indexing

Make the public page and its essential CSS, JavaScript, and images accessible to Googlebot. Then check the content that can be indexed.

Serve priority pages with HTTP 200. Use a 301 or 308 for permanent moves and 404 or 410 for absent pages. An empty or error page that returns 200 may still be treated as a soft 404. Review CDN and firewall rules, geographic restrictions, authentication, and cookie gates against the page and rendering resources.

Use robots.txt to manage crawling. To remove a publicly accessible URL from the index, use a readable noindex directive in the robots meta tag or an X-Robots-Tag header for resources such as PDFs. Google needs access to see the directive. Use authentication and access permissions for private material. [3] [4] [5] [12]

What to verify

- Verify response status, redirect destinations, robots.txt, the robots meta tag, and X-Robots-Tag from a public network.
- Compare URL Inspection’s indexed record and live test. A successful live test shows that the page can be accessed at the time of testing. Use the indexed record to check indexing.
- Test a live page, a permanent redirect, and a retired URL. Keep the expected and observed responses.

Snippet controls and Search generative AI settings

Decide what must be restricted, then choose the appropriate property setting, preview directive, or access control.

The Search generative AI setting in Search Console controls whether a property participates in applicable AI features, including AI Overviews, AI Mode, and generative AI features in Discover. Exclusion prevents the site's links and content from being shown to users or grounding responses in those experiences. It can remain eligible for ordinary Google Search, and the setting does not govern model training. Check whether the property inherits its setting.

The nosnippet directive restricts text snippets and prevents the page content from being used as a direct input for AI Overviews and AI Mode. The max-snippet directive limits text preview length. Use data-nosnippet on selected HTML sections and noindex for indexing. Google-Extended is a separate control with a documented scope. [1] [5] [6] [13]

What to verify

- Record the intended policy, owner, property scope, and effective date. Verify the response and rendered directives for new pages.
- After changing the policy, check the resulting state and allow time for recrawling, settings to take effect, and caches to update.

Rendering and accessible buying evidence

Present specifications, materials, units, price terms, revisions, and contact details as accessible page text.

Google can render JavaScript. For key pages, prefer a reliable implementation such as server rendering or static prerendering, then verify the rendered mobile output. Avoid placing critical evidence solely in images, click-only widgets, authenticated APIs, or third-party scripts. Add a text summary and a readable table to an interactive chart.

Use meaningful headings, table headers, and link text. Keep a minimum order quantity, lead time, or price range next to its unit, market, testing method, or quotation conditions. Show the technical file, sample record, or source dataset that substantiates the claim, and date the revision. [1] [2] [7] [12]

What to verify

- Compare the initial HTML with the URL Inspection render. Verify that the primary text, tables, and links are available.
- On mobile, a slow connection, and after an interaction error, check that the key facts and inquiry path remain available.

URL architecture and multilingual coverage

Give each indexable URL a distinct buying task. Connect related pages through internal links, canonicals, sitemaps, and hreflang.

Use real <a href> links between products, industry pages, comparisons, technical evidence, and inquiries. Put the intended canonical, indexable URLs in the XML sitemap. Reserve lastmod for meaningful changes. Align internal links, redirects, sitemap entries, and declared canonicals on the preferred URL. Check the canonical Google selected in Search Console.

Use separate URLs for Chinese, English, and Spanish. Configure reciprocal hreflang with valid language or language-region codes, including the page itself and relevant alternates. Localize the main text, title, navigation, and inquiry path. A distinct, indexable translation normally has its own self-referencing canonical. [8] [9] [10] [14]

What to verify

- Follow the links to evidence, each language switch, and the inquiry flow from priority entry pages. Confirm the intended 200 destinations.
- For each alternate set, verify reciprocity, language codes, canonicals, response codes, and sitemap agreement.

Structured data and visible facts

Use JSON-LD to describe the content already on the page. Choose applicable types and keep marked-up claims consistent with what the visitor sees.

Use Product, Article, Organization, or Dataset where the page actually represents that type. Reuse stable identifiers for genuine relationships. Avoid inventing a price, stock status, review, or discount for an item available by quotation. If the page qualifies a commercial fact, carry that qualification through the data.

Generate the visible specification table and JSON-LD from the same maintained source where possible. Explain what a linked attachment can substantiate. For a public dataset, state sample coverage, collection or update date, units, calculation rules, and how to obtain the source data. [1] [2] [11]

What to verify

- Check supported features with Rich Results Test. Inspect other fields in source or with a suitable Schema validator.
- Sample entities, prices, currency, availability, bylines, dates, and download links against the visible page. Validation does not guarantee a rich result or AI exposure.

Release acceptance and regression checks

Define a repeatable release checklist for each template. Fix failures that block crawling, indexing, or inquiries first.

Maintain an audit sample organized by page type, language, and purpose. Recheck the same production URLs for responses, canonicals, mobile rendering, internal links, markup, preview directives, and forms. Record the deployment date, owner, reviewer, findings, and rollback path.

Start with a small set of representative pages before rolling out a template. After deployment, check crawling, Google-selected canonicals, indexing coverage, and key entry pages. On a migration, verify old URL redirects and the new sitemap. Version claims involving price, compliance, and lead times, with effective dates. [3] [5] [7] [8] [9] [12]

What to verify

- Save pre- and post-release results for the same URL sample. Prioritize 4xx/5xx errors, invalid canonicals, accidental noindex, and broken inquiry forms.
- Record site changes, crawl status, and retest dates. Use these records to investigate seasonal or platform changes before attributing a movement to the release.

Impressions, citations, and sales outcomes

Maintain distinct records for platform reporting, direct observation, and on-site behavior. Use the right denominator and observation window for each.

The Search Console generative AI performance report currently aggregates organic link impressions across AI Overviews and AI Mode. It can be analyzed by page, country, date, and device. It does not provide a dedicated AI click count, CTR, or a split between the two surfaces. Search Labs experiments are excluded. If the report is absent, check the property's settings and consider insufficient impressions.

Use a stable set of buyer questions. Record the surface, market, language, device, observation time, answer, and visible sources. Use GA4 for landing-page behavior and conversion, and the CRM for verified inquiries and quotations. Mark uncertain acquisition sources as unknown. Ordinary Web clicks divided by AI impressions do not give an AI CTR. [1] [6] [15] [16]

What to verify

- Retain numerator, denominator, sample size, missing observations, and version. Search Console dates follow Pacific Time.
- Compare trends by page, market, and week. Keep a baseline and a comparable group before making causal claims.

Sources

[01] AI features and your website

[02] Optimizing for generative AI features in Google Search

[03] Google Search technical requirements

[04] Introduction to robots.txt

[05] Robots meta tag, data-nosnippet, and X-Robots-Tag

[06] Search generative AI control

[07] JavaScript SEO basics

[08] Canonicalization

[09] Build and submit a sitemap

[10] Localized page versions

[11] General structured data guidelines

[12] URL Inspection tool

[13] Google crawlers and Google-Extended

[14] Make links crawlable

[15] Generative AI performance report

[16] Using Search Console and Google Analytics data for SEO