Answerova

Google GEO Research

Google GEO White Paper for B2B Exporters

A practical whitepaper on market discovery, supplier evidence, and qualified inquiries

November 2026

Read online or switch languages

Contents

A practical whitepaper on market discovery, supplier evidence, and qualified inquiries

03 How AI search fits the purchasing journey

04 Map pages to buyer decisions

05 Turn a quotation page into verifiable evidence

06 Localize market evidence and the inquiry path

07 What to publish in four export sectors

08 Evidence for a manufacturing cluster

09 Keep product facts consistent across outputs

10 Six files and the teams that maintain them

11 Publish benchmarks and comparable scenarios

12 Carry buying context into a qualified inquiry

13 Technical checks before and after publishing

14 Design separate Google and Meta studies

15 A ninety-day implementation plan

16 Sources

How AI search fits the purchasing journey

Plan around the decisions international buyers make before they shortlist a supplier, request samples, or ask for a written quote.

Generative engine optimization (GEO) studies whether a source is visible in generated answers and whether those answers represent its evidence accurately. In Google Search, AI Overviews can appear within results and AI Mode offers a conversational surface. Google says these features can issue several related searches across subtopics. An importer's question may therefore lead to separate searches for product specifications, local conformity, lead time, and landed-cost assumptions. [1][2][16]

Build a question inventory from actual calls, quote requests, search data, and messages that sales can share with permission. Give each question a buyer role, destination, product family, and decision. "Can this lighting model be specified for a Saudi hotel?" and "What is the minimum order for a UAE distributor to run a trial?" require different evidence and different next steps.

Google describes indexing and eligibility to appear with a snippet as prerequisites for supporting links. Teams can improve the underlying page by making original evidence easy to find, read, and verify. Record observed presence or absence separately by surface, date, market, and language. Eligibility alone provides no guarantee that a page will be shown or cited. [1][2]

What to verify

- Collect questions that buyers have actually asked and retain their wording where possible.
- Assign each question a source of evidence and the next buyer action.
- Keep platform observations separate from aggregate search performance.

Map pages to buyer decisions

Organize the site around destination, product category, buyer role, and purchasing stage.

A freight page, a country import guide, and a procurement tutorial have separate jobs. The freight page explains route, timing, and handoff responsibilities. The country guide explains which product and destination rules buyers must verify. The tutorial walks a first-time importer through sourcing, inspection, consolidation, and delivery. A category page establishes what the supplier actually makes and which manufacturing routes it supports.

Use titles that state the decision, such as “Importing display fixtures into the UAE” or “Prototype or production tooling for molded packaging.” Pages that promise the same outcome and rely on the same evidence are candidates for consolidation. If regional rules, buyer constraints, or available products materially differ, make that difference clear in the opening section and in the internal links. [2]

Connect each explanatory page to one practical next action. A standards page can link to a model-specific report. A cost page can point to its sample dataset and an inquiry form that carries the destination. A comparison can prefill the two options and the quantity being evaluated. Keep links in the main copy where the buyer needs them.

What to verify

- Assign each important URL one main decision and one practical next action.
- Use descriptive links from related products, destinations, and guides.
- Consolidate overlapping pages after checking intent, evidence, and the action offered.

Turn a quotation page into verifiable evidence

A buyer can assess a supplier faster when the main claim, its conditions, and the documents that support it appear together.

A public cosmetics quotation guide illustrates the structure. It opens with buyer and product context, explains the manufacturing and labeling options, separates initial and recurring costs, and closes with a checklist for requesting quotations. The underlying fields are useful across sectors. A buyer comparing custom textiles can apply the same structure to fabric weight, sizes, decoration, setup charges, packout, and repeat-order terms. [5]

Attach each numerical claim to its scope. An MOQ can apply to a single color and size, an approved stock formulation, or a custom mold. A lead time can begin after artwork approval, payment, or component availability. If the visible page gives a planning range, show what can move the buyer outside it and what must be confirmed in writing.

Publish photos and diagrams that answer specific questions. Caption a sample with the material, finish, and manufacturing stage it shows. Mark inspection points, acceptance criteria, and who is responsible at each handoff. Link to a specification, sample sheet, test method, or authority where the buyer can verify the claim. Label a specimen as a specimen and redact customer identifiers before publication.

What to verify

- Place the short answer and important conditions above the detailed explanation.
- Pair prices, MOQs, and lead times with a unit, scope, and review date.
- Give purchasers a reusable list of fields for their own RFQ.

Localize market evidence and the inquiry path

A translated page is useful when the destination assumptions, business terms, and next action also work for local buyers.

Use a stable, crawlable URL for each supported language. Publish a complete main text, title, navigation path, and inquiry route in that language. For Arabic, test right-to-left layout alongside numbers, model IDs, and dimensions. A language switch helps people move between editions. An indexable locale URL lets search engines discover each edition directly. [3]

Where equivalent pages exist, connect all language versions through reciprocal hreflang, including self-references. Keep canonicals consistent with the pages intended for indexing. Use separate country editions when conformity, shipping responsibility, payment terms, currencies, or product scope changes the decision. Where no regional difference affects the answer, one clear edition in that language can serve several markets. [3]

Test the full purchasing journey from the localized landing page to quote confirmation. Check that links to calculators, downloadable specifications, forms, and support maintain the expected language or clearly describe a change. Record whether prices are estimates or formal quotations. Verify who responds, the channel offered, and which commercial documents are available in the buyer’s language.

What to verify

- Test reciprocal hreflang, canonical destinations, and live HTTP responses.
- Check the display of prices and dimensions on mobile in each edition.
- Submit a test inquiry and verify its destination, fields, and confirmation language.

What to publish in four export sectors

The same question map yields different evidence for a garment, a package, a luminaire, and a furniture project.

Apparel and textiles. Publish fiber composition, fabric weight, size and color tolerances, sampling steps, and MOQ by color or size. Identify the records needed for labeling and origin. Separate prototype approval, production lead time, and repeat orders. For custom packaging, document dieline, substrate, print method, finish, and tooling charges. Explain how sample approval, color differences, and the MOQ for each process affect the quote.

Lighting. Connect each model to installation conditions, electrical specifications, test evidence, and available spares. State the model, market, test date, and scope that a report covers. A project buyer needs to know which documents the manufacturer can supply and which conformity steps the importer must verify for that destination.

Furniture and project fit-out. Provide a bill of quantities template, material and finish options, dimensional tolerances, sample approval steps, and quality checks by production lot. Include packing volume and delivery milestones. Across all four sectors, collect the evidence during quoting and quality review, while the source and scope are easy to confirm. Give technical and sales teams responsibility for reviewing the facts before release.

What to verify

- Start with the three purchasing questions that repeatedly delay a quotation.
- Tie each model or capability to evidence with a known owner.
- State which requirements apply at sample, production, and delivery stages.

Evidence for a manufacturing cluster

Show what happens between choosing a material and clearing the finished order for shipment.

An archived guide to the Chaoyang sleepwear cluster pairs four fact cards with five purchasing pitfalls. The cards cover the cluster's products, its typical buyers, the Gulf retail season, and labeling or conformity questions to verify for the destination. The pitfalls are a bulk fabric lighter than specified, shade variation across dye lots, shrinkage after washing, separate factory minimums for private-label components, and a production schedule that leaves too little time for approval and transit. Each problem points to a specific decision the buyer can make before ordering. [17]

Turn the guide into a sampling workflow. Choose the material, shade, and hand feel from physical swatches. Approve the fiber and care label for the destination. Keep the signed specification, retained sample, weight tolerance, and shrinkage limit. Confirm the minimums for woven labels, care labels, tags, and packaging before promising a launch. Map approval, production, inspection, packing, paperwork, and transit back from the required shelf date. Mark who signs each approval and what the buyer should include when sending an RFQ. [17]

Use visual evidence where it answers a question. Caption the color-card wall and care-label specimen with what each image actually shows. Give process diagrams text labels and a readable description. Publish destination or product comparisons as HTML tables with units and an effective date. Link to the CSV behind any benchmark, and expose the exchange rate, weight, freight, and tax assumptions used by a calculator. The archived 150-row sourcing dataset illustrates how readers can reproduce the input statistics. Keep manufacturer claims, dated planning assumptions, and formal quotations distinguishable. [6][7][17]

What to verify

- Show the evidence, tolerance, and approval owner for each sample decision.
- Connect the photo, table, and download to a question a buyer can act on.
- Verify the current conformity rules for the exact product and destination before quoting.

Keep product facts consistent across outputs

A shared set of verified business facts makes the catalog, specifications, markup, calculator, and form easier to maintain.

Separate the product catalog, pricing parameters, inquiry rules, and publication logs. Catalog entries should have stable identifiers and explicit units. The price configuration should distinguish currency conversion, freight, tax assumptions, and their effective dates. Inquiry rules describe the fields and routing logic. Use stable identifiers in question sets and append-only observation logs to compare research rounds. Keep private inquiry and operations data outside public exports.

Generate visible specifications and structured data from the same verified source where feasible. Link organizations, pages, products, offers, and downloadable datasets with stable identifiers. If a catalog showcases manufacturing capabilities or representative samples, its markup should reflect that status. Verify that every public figure and availability statement matches the page content. Google’s retirement of FAQ rich results leaves useful buyer questions in the visible copy. [4][12][14]

Check the live HTML as well as the rendered page. Review the first paragraph for a direct answer, then verify the substance and provenance of the figures. Numeric density can flag material for review, but it does not establish quality or eligibility. Keep receipts for changed URLs and run a separate check for indexing. Google documents ordinary search eligibility without a requirement for a special AI file or markup. [1][2]

What to verify

- Regress one product, price, or lead-time change across all public outputs.
- Keep placeholders, historic samples, and live offers distinguishable.
- Give each data source a maintainer and an effective date.

Six files and the teams that maintain them

Separate the inputs that define the offer from the records used to publish, test, and measure it.

In a fixed public implementation, pricing.json holds dated rates and market assumptions, with 7.15 meaning Chinese yuan per US dollar. products.clean.json defines nine categories and 150 examples across seven of them, with explicit fields for FOB, weight, MOQ, and lead time. order-intent.json describes service scope, validation, and delivery channels. Two of its eight inquiry fields are mandatory. Separate ownership works well here. Product staff maintain specifications, quoting staff approve rate inputs, and sales owns required fields and routing. [13][20]

Three further files serve operations. indexing_queue.json keeps pending, completed, and content-hash entries. The archived inspection found 0, 133, and 131 entries respectively. Completed records acknowledge a submission response. Search indexing is verified independently. cite_monitor_queries.json preserves 16 stable bilingual question IDs. cite_monitor_log.jsonl contains 32 append-only observation records from two historical checks. Its citation flag included some search-result appearances, so use separate labels for a results listing and an actual linked answer. [19][20]

Connect the visible page to its underlying entities. Link an ItemList to products and offers, a Dataset to its downloadable CSV, and a service to its provider with stable identifiers. An archived review of 121 HTML pages found 137 JSON-LD blocks and 70 pages with 105 visible tables. It also found English catalog markup for 150 products where 14 cards were visible, and sample products flagged as in stock. Treat counts as observations from that revision. Compare offers and availability against the rendered pages in each language. Keep internal queues and buyer data private unless a specific public purpose justifies publication. [4][14][20]

What to verify

- Document an owner, version, and unit for every business input.
- Check that structured offers, availability, and downloadable data match the visible page.
- Separate submission acknowledgments, indexing evidence, and answer citations.

Publish benchmarks and comparable scenarios

A procurement benchmark becomes useful when a reader can reproduce the result or replace the assumptions with their own.

A public China-to-Gulf sourcing benchmark uses 150 catalog samples across seven categories. Its CSV reports observations such as FOB unit prices, MOQs, weights, and catalog lead times. The benchmark uses category-level inputs and destination assumptions to model landed cost. For home fragrance, the historical model used 22 samples, a median FOB of US$1.76, a median MOQ of 48, and a median weight of 0.30 kg per unit. Its estimated per-unit totals were about US$4.20 for the UAE and US$4.60 for Saudi Arabia. These are figures from that published model. [6][7]

For a comparison, give both routes the same basket and destination. A published scenario models a US$3,000 order to Riyadh at approximately US$4,290 for platform sourcing and US$3,600 for a sourcing agent. The difference depends on the author's freight, service, and handling assumptions. The example does not isolate taxes or buyer time. A real comparison needs the same specification, quantities, trade terms, and verified charge scope. [8]

Show the sample date and size, then link the data and describe the formula and any exclusions. Mark catalog observations, modeled estimates, and binding quotes separately. Run identical inputs through the main page and calculator after each rate update, then investigate differences beyond rounding. [6][7][14]

What to verify

- Publish sample scope, collection date, and definitions beside the table.
- Use the same inputs and clearly defined cost coverage for both routes.
- Keep source data downloadable and date the assumptions used in the model.

Carry buying context into a qualified inquiry

Keep the first step easy and let sales collect the fields that make a written quotation possible.

Pass destination, product or comparison URL, selected model, and language into the inquiry flow. A product page can send the MOQ and product reference. A country page can pass the import destination. A sourcing comparison can pass both routes. A public sourcing implementation starts by requesting destination and a product reference, then adds optional detail on category, quantity, private label, budget, timing, and notes. [13]

Track four separate milestones. An inquiry click opens a channel. A received inquiry enters the inbox or CRM. A qualified inquiry has a reachable contact, destination, quantities, specification, and timing that sales can work with. A written quote has a version, currency, expiry, and price scope. In WhatsApp and email, verify actual delivery or receipt rather than counting a prefilled message as sent.

Join page and channel data to the CRM with an inquiry identifier. Deduplicate before reporting conversion, store buyer-reported source separately from confirmed tracking, and preserve an unknown category. Give sales a way to flag questions the page did not answer. Use those questions to refine the product evidence and make the handoff more useful.

What to verify

- Carry page and locale context without putting personal data in public URLs.
- Verify that one test lead moves through all four milestones.
- Report channel clicks, qualified inquiries, and quotations separately.

Technical checks before and after publishing

Use automated flags to find pages that need review. Let editors and commercial teams verify the underlying facts.

Inspect both the original HTML and the rendered page. A historical audit script stripped markup and scripts, flagged English pages under 120 visible words, and warned at 120–249 words. It exposed the first 200 visible characters for a human to check whether the opening answers the purchasing question. An internal editorial template sought at least three numeric signals per 100 English words. The counter can pick up years and header text. Apply judgment by language and verify the source, units, date, and applicability of every meaningful number. Google sets no such word or number threshold for AI features. [1][2][18][20]

Check titles, viewport metadata, HTTPS canonicals, language references, robots directives, and index intent. Verify primary content, quote terms, navigation, and JSON-LD on a live URL for each locale. Test an identical basket in the product flow and calculator after rates change. A historical correction found an FAQ stating an MOQ of 30–200 for home fragrance where the underlying sample data showed 12–300. Fix the data source and regenerate all affected outputs. [3][4][14]

Preserve an immutable copy of the article, input data, and build revision used for the audit. Compare content hashes for changed HTML and public assets such as the sitemap, robots file, and inquiry contract. One November 2026 release logged 102 HTML submissions and five additional non-HTML submissions after a gap was found in the deploy script. Keep the submitted URL list and response status, then record crawl and indexing checks separately. If AI answer visibility matters, store that observation with its own question, market, date, and evidence. [18][19]

What to verify

- Save a dated release snapshot with its article, data, and calculation versions.
- Check rendered content and quotations from every supported language.
- Keep editor review, deployment success, indexing, and answer visibility as distinct outcomes.

Design separate Google and Meta studies

This is a proposed measurement design. It provides no observed results for Google AI Overviews, AI Mode, Meta AI, or Muse.

Create a fixed set of non-brand questions and a separate brand set. One design uses four industries, two markets, two languages, five buying stages, and three questions per combination, for 240 questions. Running both Google surfaces on three observation days creates up to 1,440 scheduled checks. Save the exact prompt, date and time, market, language, device, and answer or screenshot. Record whether a feature was available, whether it appeared, and which domains or URLs were visible.

Use distinct denominators for answer appearance, brand mentions, linked citations, and citation accuracy. Audit a sample of cited passages against the source page. Search Console's generative AI report shows impressions and can be examined by page, country, device, and date. Compare it with overall search performance and site analytics. An impression does not establish a visit, a qualified lead, or the accuracy of an answer. [9][15]

For Meta, select a matched subset, for example 32 questions. Record the precise surface tested, such as a Meta AI entry point or the separate Muse personal agent where available, and whether the interface identifies the underlying model. Analyze this comparison separately from Google and exclude unavailable runs from paired comparisons. Muse Spark is a model. Meta introduced Muse separately as a personal agent. [10][11]

What to verify

- Agree on source, citation, and accuracy definitions before sampling.
- Keep Google surfaces, Meta products, and unavailable runs distinguishable.
- Report denominators and preserve evidence for independent review.

A ninety-day implementation plan

Start with an export segment that already has credible technical evidence and a team ready to handle inquiries.

Days 1–14. Inventory URLs, markets, and language paths. Interview sales and pick one industry and two destination markets. Write the initial question map, gather confirmed product evidence, and record the first visibility baseline. Assign one owner each for content, technical release, and quote follow-up.

Days 15–45. Build or revise the product, country, comparison, and quote pages for the priority decisions. Connect their supporting files and localized forms. Check live response, canonical and language relationships, visible content, structured data, and quote routing. Exercise at least ten test cases covering typical, incomplete, and conflicting inputs. Set any internal target for click depth and contextual links before auditing it.

Days 46–90. Repeat the visibility observations on the same questions and conditions where possible. Inspect whether cited passages accurately describe the offer. Compare inquiry and quote quality with the baseline and an appropriate comparison set, when available. Document platform, seasonality, and rollout changes before interpreting differences. Prioritize repairs where there is demand and the missing evidence can be supplied.

What to verify

- Confirm a live, localized path from a product or country page to a received RFQ.
- Release only claims backed by an owner, scope, source, and review date.
- Deliver a measurement record that another analyst can repeat.

Sources

[01] Google Search Central · AI features and your website

[02] Google Search Central · Optimizing for generative AI features

[03] Google Search Central · Localized versions of your pages

[04] Google Search Central · General structured data guidelines

[05] UniBeauty · Private-label cosmetics quotation guide

[06] SourceToGulf · China-to-Gulf sourcing benchmark

[07] SourceToGulf · Benchmark CSV

[08] SourceToGulf · Published comparison model, fixed revision

[09] Search Console Help · Generative AI performance report

[10] Meta · Introducing Muse Spark

[11] Meta · Introducing the Muse personal agent

[12] Google Search Central · FAQ rich-result deprecation

[13] SourceToGulf · Public inquiry fields, fixed revision

[14] SourceToGulf · Technical audit and pricing corrections

[15] Google Search Central · Using Search Console and Google Analytics together

[16] Aggarwal et al. · Generative Engine Optimization, research paper

[17] Industrial-belt page, fixed revision

[18] Technical audit script, fixed revision

[19] November release log, fixed revision

[20] Source and data files, fixed revision