

THE INDEX CLUB / FIELD MANUAL 02

The Ecommerce Search Playbook.

From catalogue to customer: a practical operating system for making products findable, understandable, recommendable and buyable across Google, Shopping and AI.

Why this exists.

Ecommerce SEO has a strange problem: almost every tactic sounds sensible in isolation. Optimise product pages. Add category copy. Fix faceted navigation. Improve schema. Submit a feed. Build links. Publish buying guides. The difficulty is not knowing that these things exist; it is deciding which parts of the catalogue deserve to exist, which system is actually constraining growth, and which intervention will create commercial leverage rather than another completed task.

This playbook treats ecommerce search as one connected system. A product can fail before ranking is even relevant: the URL may be unreachable; the catalogue model may be incoherent; the product identity may be incomplete; the query may map to the wrong page type; the feed may disagree with the PDP; the product may be surfaced but unpersuasive; or measurement may lose the sale. Google itself separates website crawling, structured data and Merchant Center feeds as complementary ways of understanding products. [\[G1\]](#)[\[G3\]](#)[\[M4\]](#)

THE CENTRAL IDEA

Do not optimise pages in isolation. Optimise the catalogue-to-customer system: eligibility → understanding → discovery → competition → selection → persuasion → measurement.

The evidence layer prioritises platform documentation, standards and large-scale UX research. The practitioner layer adds Warren Hance's published work, working methods and Delphi-trained point of view. The Index Club layer turns both into a process a team can actually run. Opinion is labelled as opinion; platform behaviour is not inflated into a universal ranking factor.

LAYER	WHAT IT CONTRIBUTES
PLATFORM / STANDARD	What Google, Merchant Center, OpenAI, Shopify, Schema.org and GS1 explicitly support or require.
RESEARCH	What usability or large datasets suggest about shopper behaviour and implementation quality.
WARREN'S TAKE	The practitioner judgement layer: how to prioritise, where common advice breaks down, and what matters commercially.
INDEX CLUB PROCESS	A repeatable implementation or diagnostic that can be applied to a real store.

Contents.

SECTION	THE JOB
01	The ecommerce search stack
02	Catalogue-to-Customer
03	Demand & catalogue architecture
04	Technical eligibility
05	Facets, filters & pagination
06	Category & PLP systems
07	Product detail pages
08	Product data, feeds & schema
09	Inventory & lifecycle
10	Content, internal search & authority
11	AI shopping & agentic commerce
12	International, migrations & scale
13	Measurement & prioritisation
14	The 90-day operating plan
15	The 150-point audit
16	Myth book
17	Source library

PART 01

The ecommerce search stack

Search did not become one channel. It became a stack of discovery systems sharing the same product truth.

A product can be discovered before a classic blue link ever wins.

Google can understand ecommerce products through crawlable pages, structured data and Merchant Center feeds; richer merchant listings can expose price, availability, shipping and returns, while Merchant Center is required for some shopping surfaces. [\[G3\]](#)[\[G4\]](#)

In parallel, OpenAI now publishes a product-feed specification for product discovery in ChatGPT, and Shopify's agentic storefront documentation describes product distribution to AI channels including ChatGPT, Google AI Mode/ Gemini, Microsoft Copilot and Meta for eligible merchants. [\[O1\]](#)[\[S6\]](#)[\[S7\]](#)

THE IMPORTANT BIT

This does *not* mean “SEO has been replaced by feeds”. It means the same product can now be understood through multiple representations. The job is to keep those representations aligned.

The modern ecommerce discovery stack

LAYER	WHAT THE CUSTOMER MAY SEE	WHAT THE SYSTEM NEEDS
Web search	Category results, PDPs, images, editorial results	Crawlable pages, intent match, authority, internal structure.
Google shopping surfaces	Free listings, merchant results, product knowledge experiences	Accurate Merchant Center data, offers, identifiers, policy data.
AI discovery	Synthesised recommendations and product selections	Relevant product facts, feeds where supported, credible evidence and retrievable web content.
On-site discovery	Search, categories, filters, recommendations	Clean taxonomy, attributes, inventory and useful UX.
Conversion	PDP, cart, checkout	Trust, price, availability, delivery, returns and a persuasive decision experience.

WARREN'S TAKE

Traditional organic search and Google Shopping still deserve first-order attention for most ecommerce brands. AI shopping is important, but commercial adoption and measurement remain uneven. The sensible move is not to build a parallel ‘AI SEO’ programme; it is to make the same catalogue cleaner, richer and more trustworthy so established and emerging surfaces can use it. [\[W2\]](#)[\[W3\]](#)[\[W4\]](#)

Google's own ecommerce guidance still centres accessible architecture, product data, structured data, review quality, URLs and pagination. The future-facing work is additive: richer product attributes, better feed consistency and machine-operable commerce become extensions of the same foundation. [\[G1\]](#)[\[G2\]](#)[\[G6\]](#)[\[M2\]](#)[\[M6\]](#)

PART 02

Catalogue-to- Customer

Seven gates. Diagnose the failing gate before prescribing work.

The Index Club Catalogue-to-Customer model.

Most ecommerce audits are organised by discipline: technical, content, links, schema. That is convenient for specialists but not for diagnosis. Catalogue-to-Customer is organised around the sequence a product must survive before search visibility becomes revenue.

01	ELIGIBLE	Can the system access the page, product or feed record?	HTTP, robots, canonical, links, sitemaps, feed eligibility.
02	UNDERSTOOD	Does the system know what the product is and how variants/offers relate?	GTIN, brand, variant group, taxonomy, structured attributes.
03	DISCOVERABLE	Does a durable page or record match the way customers search and browse?	Demand mapping, categories, facets, internal search, intent.
04	COMPETITIVE	Does this result give more useful evidence than equivalent inventory elsewhere?	Selection, expertise, reviews, original content, authority.
05	SELECTED	Does the product/page actually earn visibility in Search, Shopping or AI discovery?	Rankings, listings, recommendations, retrieval, share of voice.
06	PERSUASIVE	Once discovered, does the experience help a customer choose and buy?	Price, delivery, returns, imagery, comparisons, trust, UX.
07	MEASURED	Can the team connect search discovery to product and revenue outcomes?	Revenue, CVR, AOV, margin, attribution, platform reporting.

THE INDEX CLUB PROCESS

For any underperforming category or product set, start by naming the failing gate. A collection that is not indexed has an eligibility problem. A feed with missing GTINs has an understanding problem. Two overlapping collections have a discovery/architecture problem. A ranking PDP with poor conversion has a persuasion problem. This prevents teams from prescribing 'more content' to every symptom.

A useful audit question

If we fixed this issue perfectly, **which gate would improve and what commercial mechanism would change?** If the team cannot answer that, the task probably needs reframing.

PART 03

Demand & catalogue architecture

A catalogue is an information architecture for products, not a folder tree copied from the ERP.

The right page begins with the right job.

Shopify's 2026 SEO guidance gives a useful starting distinction: product pages suit specific product/model/use-case searches, collection pages suit broader comparison intent, and guides suit educational tasks. [\[S1\]](#)

That sounds elementary until a catalogue has five thousand products, campaign landing pages, tags, brand hubs, search pages and filters all competing for similar language. The strategic work is deciding which stable page should own the intent.

QUERY / TASK	LIKELY BEST OWNER	WHY
"women's running shoes"	Collection / PLP	The shopper wants to compare options.
"ASICS Gel-Kayano 32 black size 10"	PDP / selected variant	A specific purchasable item is implied.
"waterproof hiking boots wide fit"	Curated subcategory or durable facet page	Commercial comparison with durable attributes.
"trail running shoes vs road running shoes"	Guide / comparison	The customer needs criteria before choosing a category.
"best protein powder for baking"	Evidence-led guide with product routes	Use-case selection may require expertise before purchase.

The four tests before creating another collection

- **Demand:** Is there a real customer task or recurring search pattern?
- **Inventory:** Is there enough suitable stock to satisfy the task?
- **Differentiation:** Does this page show a meaningfully distinct set, rather than a renamed parent category?
- **Durability:** Will the concept survive merchandising changes and still deserve a permanent URL?

WARREN'S TAKE

Adding more collection pages can be one of the most powerful levers in ecommerce SEO - but only when the page represents a real, durable way customers want to shop. Programmatic scale should make good decisions repeatable; it should not industrialise thin pages. Automation clears the repetitive work so human judgement can stay focused on intent, commercial value and exceptions. [\[W5\]](#)[\[S2\]](#)

THE INDEX CLUB PROCESS

Build a query-to-page map using Search Console, keyword data, internal search, paid search terms and catalogue attributes. Assign one preferred URL to each meaningful intent cluster. Flag: missing pages, wrong page type, cannibalisation, thin inventories and high-paid-spend queries with weak organic ownership.

Prioritise commercial architecture before informational volume.

For most retailers, the biggest organic upside is often buried in money pages that already have inventory, conversion data and paid demand. A strong roadmap therefore begins with collection/PDP coverage and technical accessibility, then uses editorial content to solve decisions and strengthen those commercial hubs - not to hit a monthly article quota.

PART 04

Technical eligibility

The crawler cannot buy what it cannot reach. The feed cannot rescue a catalogue whose page states disagree.

Crawlability is a catalogue property.

Google explicitly recommends crawlable navigation from menus to categories, subcategories and products. It also notes that products reachable only through an internal search box may not be discovered by crawling, because Googlebot generally does not submit site-search forms. [\[G2\]](#)

For large catalogues, technical SEO is less about finding isolated errors and more about controlling the URL system: parameters, facets, sorting, pagination, variants, campaign URLs, internal search and tracking can multiply crawlable states far beyond the number of products actually sold.

The technical baseline

- Stable 200 responses for pages that should remain live.
- Consistent canonicals to the preferred indexable URL.
- Crawlable `<a href>` internal links for important destinations. [\[G13\]](#)
- XML sitemaps containing canonical, indexable URLs only.
- Pagination or equivalent crawl paths that expose products beyond the first loaded set. [\[G7\]](#)
- No accidental noindex/robots/WAF rules on commercial templates.
- Rendered HTML containing the information customers and search systems actually need.
- Parameter policies designed before the catalogue creates millions of combinations.

WHAT THE EVIDENCE SAYS

Google's crawl-budget guidance is aimed mainly at very large or rapidly changing sites. Smaller stores should not manufacture a crawl-budget project when Search Console and fresh sitemaps show normal discovery. [\[G14\]](#)

Infinite scroll is a UX pattern, not a crawl strategy.

Google says its crawlers generally follow URLs found in href attributes and do not click buttons or trigger user actions to reveal more content. If a PLP uses load-more or infinite scroll, provide crawlable URLs or another reliable discovery path for the full product set. [\[G7\]](#)

WARREN'S TAKE

The question is never 'does this framework use JavaScript?' The question is whether the final product/category state is reliably renderable, linked and discoverable at the scale the business operates. Technical SEO should model the failure mode, not repeat a platform stereotype.

PART 05

Facets, filters & pagination

Filters are essential shopping UX. They are not automatically pages Google should index.

Do not confuse a filter state with a search landing page.

Faceted navigation helps shoppers narrow large product sets, but every filter combination can create another URL state. Shopify and Ahrefs both document the familiar SEO risk: redundant combinations, crawl waste and index bloat can grow much faster than the useful catalogue itself. [\[S3\]](#)[\[H1\]](#)

At the same time, the underlying attributes can reveal valuable long-tail demand. The answer is neither 'index every facet' nor 'block every facet'. It is to distinguish interaction states from durable landing pages.

The Index Club facet decision framework

QUESTION	NO	YES
Is there meaningful recurring demand for exactly this filtered product set?	Keep it a shopping control.	Continue.
Is the product set materially different from the parent page?	Serve the need on the parent or another existing page.	Continue.
Is the concept durable across stock and seasonal changes?	Do not create a permanent search landing page.	Continue.
Can the store keep enough relevant inventory to satisfy the task?	Do not index a fragile/empty experience.	Continue.
Would a curated URL, copy and links be better than a raw parameter?	Usually build a deliberate landing page.	Only then consider the facet URL itself as the destination.

THE INDEX CLUB PROCESS

For every facet, create one policy row: facet name → customer purpose → search demand → URL pattern → crawl rule → index rule → canonical rule → promoted landing pages → review date. The policy should exist before developers or apps decide the behaviour by default.

UX and SEO should share the same attribute model.

Baymard's research repeatedly finds that shoppers need useful category-specific filters and clear visibility of applied filters. In 2026 it reported widespread PLP usability problems, including missing applied-filter overviews and unclear filtering. [\[B1\]](#)[\[B2\]](#)

This is a useful discipline for SEO too: if an attribute is important enough for customers to refine by, it is worth asking whether it also belongs in product data, comparison content, internal search synonyms or a durable landing page.

WARREN'S TAKE

I would rather build ten intentional collection pages around validated product attributes than expose ten thousand parameter combinations and hope Google finds the good ones. Demand, inventory, differentiation and permanence come before indexation.

PART 06

Category & PLP systems

The category page is where search intent, merchandising and product discovery collide.

Treat PLPs as search-aware merchandising systems.

Category/collection pages often represent the broad commercial jobs individual PDPs cannot satisfy: compare brands, styles, use cases and price points. Google recommends a crawlable hierarchy through category and subcategory pages; current Shopify guidance likewise frames collections around broader searches where shoppers want options.

[G2][S1]

Above the fold: orient, then merchandise.

- A precise H1 that names the product set in customer language.
- A short, genuinely useful orientation - not 500 words written to satisfy an imaginary SEO quota.
- The product grid and essential comparison information visible quickly.
- Filters that match the category-specific attributes people actually use.
- Clear inventory count, sort logic and applied-filter state where useful.
- Fast mobile interaction and no intrusive content wall between intent and products.

WHAT THE EVIDENCE SAYS

Baymard's 2026 PLP research says the product list is where discovery actually happens and reports that 80% of benchmarked ecommerce sites have serious PLP usability flaws. Its wider product-list research links poor filtering/list usability with much higher abandonment in controlled testing. [B1][B3]

What should category copy actually do?

GOOD JOB

Explain selection criteria or product differences that help a shopper narrow options.

Clarify fit, compatibility, materials, use case or buying constraints.

Link to meaningful adjacent categories or expert guidance.

Demonstrate why this retailer has a credible selection or point of view.

BAD JOB

Paraphrase the H1 three times.

Stuff synonyms below the grid.

Create a wall of internal links solely to move PageRank.

Write generic manufacturer-neutral copy any store could publish.

WARREN'S TAKE

Commercial coverage comes before word count. A thin collection is not automatically one with little prose; it is one that gives the shopper and search system too little reason to understand why this set exists, how it differs, and why it is useful. Sometimes the fix is copy. Sometimes it is better products, better filters, better attributes or a better page model.

THE INDEX CLUB PROCESS

For each priority PLP, review: query intent → inventory depth → product-card information → filters → intro guidance → internal links → supporting expertise → structured breadcrumbs → organic revenue. If the page gets traffic but not sales, do not automatically add more SEO text; inspect the assortment and decision experience.

PART 07

Product detail pages

A PDP is product truth, evidence and conversion interface - not a description field with an Add to Cart button.

Make the product independently understandable.

Product pages increasingly need to work across several contexts: a human landing directly from Google; Google extracting offer details; Merchant Center verifying feed data against the page; an AI system retrieving product facts; or an agent deep-linking to a selected variant. [\[G3\]](#)[\[G4\]](#)[\[O1\]](#)

The PDP information stack

- **Identity:** clear product name, brand, model and selected variant.
- **Offer:** current price, availability, condition, promotion, delivery and return expectations.
- **Decision attributes:** dimensions, materials, compatibility, fit, ingredients, technical specifications, use cases.
- **Evidence:** original imagery, demonstrations, measurements, customer reviews and expert explanation.
- **Comparison:** why this product differs from adjacent variants or alternatives.
- **Action:** variant selection, stock state, shipping, returns and a stable route to purchase.

WHAT THE EVIDENCE SAYS

Google's Product/Offer structured-data guidance supports merchant-listing enhancements including price, availability, shipping and returns. Its variant documentation recommends ProductGroup for related variants and notes that selected variant states should be addressable and accurate. [\[G4\]](#)[\[G6\]](#)

Manufacturer copy is not the real question.

The simplistic rule 'every PDP description must be unique' often mistakes duplication for usefulness. If fifty retailers sell the same product, rewriting a manufacturer's dimensions into fifty slightly different paragraphs does not create much information gain. The stronger opportunity is to add what only the retailer can know or demonstrate: fit guidance, comparison logic, testing, compatibility, bundles, real images, customer questions, delivery constraints and experience using the product.

WARREN'S TAKE

For important PDPs, I care less about whether every sentence is unique and more about whether the page contains enough product truth and decision support to deserve the click. At scale, automate consistent baseline fields; spend human effort where judgement actually changes the customer's decision.

Images are search surfaces too.

Google recommends standard HTML image elements, discoverable image URLs, useful alt text and high-quality imagery while balancing performance. For ecommerce, additional angles and context shots also reduce uncertainty for users. [\[G9\]](#)[\[B6\]](#)

PART 08

Product data, feeds & schema

Your most important organic-search file may not be an HTML page.

Product data is content.

Merchant Center attributes are not administrative leftovers. They describe the product in structured form: identifiers, title, description, price, availability, colour, material, shipping and far more. Google says accurate required and suggested attributes help people find products, and its `product_detail` and `product_highlight` fields now explicitly mention AI-driven surfaces such as AI Mode. [\[M1\]](#)[\[M2\]](#)[\[M3\]](#)

INDEX CLUB PRINCIPLE

If a buying criterion matters to the customer, ask where it exists in four places: **PDP → structured data → merchant feed → internal catalogue model**. Missing or contradictory attributes are content defects, not merely feed defects.

Identity before optimisation

GS1 defines GTIN as a global identifier for trade items. Google says missing or incorrect GTINs can limit product visibility and warns merchants not to invent values. Stable IDs matter in OpenAI's product-feed specification too: item/variant identifiers should remain stable as price and stock change. [\[GS1\]](#)[\[M5\]](#)[\[O1\]](#)

The product truth matrix

FIELD	PAGE	SCHEMA	GOOGLE FEED	OPENAI / AGENTIC
Product / variant ID	SKU / selected state	sku / ProductGroup relations	id / item_group_id	item_id / group_id
Brand	Visible brand	brand	brand	brand
Price	Current purchasable price	Offer.price	price / sale_price	price / list price
Availability	Visible stock state	Offer.availability	availability	availability
Variant attributes	Size, colour etc.	ProductGroup / variesBy	variant fields / item group	variant_dict / group_id
Shipping / returns	Visible policy	OfferShippingDetails / MerchantReturnPolicy	shipping / return settings	shipping / returns fields

The exact field names and supported capabilities differ by system; the strategic point is consistency, not one universal schema. Google even documents an order of precedence for certain shipping/return configurations, reinforcing why teams need governance rather than duplicate unmanaged truth. [\[G4\]](#)[\[M4\]](#)[\[O1\]](#)[\[SC1\]](#)[\[SC2\]](#)[\[SC3\]](#)[\[SC4\]](#)

WARREN'S TAKE

Merchant Center should sit in the SEO conversation. If a product is disapproved, misidentified, stale or missing the attributes customers search by, that is an organic discovery problem even if the paid-media team happens to own the account.

THE INDEX CLUB PROCESS

Run a feed-to-page consistency audit on priority products: ID → title → brand → GTIN/MPN → price → sale price → availability → selected variant → image → shipping → returns → core attributes. Pull mismatches into one defect backlog owned across SEO, merchandising and feed operations.

PART 09

Inventory & product lifecycle

Stock changes. URLs should not behave as if every inventory event is the same problem.

Treat lifecycle states as different intents.

Google's ecommerce URL guidance distinguishes empty categories and recommends deliberate handling; Merchant Center likewise expects accurate availability. The broader principle is simple: temporary unavailability, permanent discontinuation and an empty category are different states and should not share one mechanical response. [\[G8\]](#) [\[M1\]](#)

Lifecycle decision matrix

STATE	DEFAULT SEARCH RESPONSE	CUSTOMER RESPONSE
Temporarily out of stock	Usually keep the durable PDP live if the product is expected back and the page still satisfies demand.	Show status, expected return if known, restock action and close alternatives.
Permanently discontinued + true successor	Consider a relevant permanent redirect when the successor genuinely satisfies the same intent.	Explain replacement/upgrade where useful.
Permanently discontinued, no equivalent	Keep useful research content if demand/equity justifies it, or retire cleanly. Avoid irrelevant blanket redirects.	Offer related categories without pretending they are the same item.
Seasonal product returning	Preserve the durable URL where appropriate.	Communicate season/status and alternatives.
Variant removed	Retain parent PDP if the product family remains valid.	Remove unavailable choice; preserve valid deep links where sensible.
Category becomes empty	Assess whether the category is temporarily empty, permanently obsolete or still useful as navigation/content.	Do not strand customers on an unexplained empty grid.

WARREN'S TAKE

The worst lifecycle rule is 'redirect every discontinued product to the category'. Relevance matters. Redirects are not a bin for expired inventory. Keep, replace or retire according to customer intent, equity and a real successor relationship.

THE INDEX CLUB PROCESS

Give every catalogue event a status code in the merchandising system: NEW / LIVE / TEMP-OOS / SEASONAL-OFF / DISCONTINUED-SUCCESSOR / DISCONTINUED-NO-SUCCESSOR. Map each status to page behaviour, feed behaviour, sitemap state, internal links and monitoring. This turns reactive SEO clean-up into governance.

Launch new products before you need them to rank.

New-product SEO is often treated as a metadata task after the PDP is live. A better launch connects the product into collections, navigation, editorial support, Merchant Center and structured data from day one, while preserving the identifiers the product will need across its lifecycle.

PART 10

Content, internal search & authority

Publish around decisions. Mine the questions and product language your customers already generate.

Your internal search box is a demand-research tool.

Customers regularly describe products differently from taxonomy teams. Internal-search queries reveal brand misspellings, use cases, attributes, problem language and combinations the catalogue may not support cleanly. Baymard notes that search is a core product-discovery path and that customers need to be able to move quickly between browse and search. [\[B5\]](#)[\[B7\]](#)

Turn internal-search data into four backlogs

- **Synonyms:** language the search engine should understand without creating a new SEO page.
- **Attributes:** recurring decision criteria that should become structured product data or filters.
- **Collections:** durable commercial demand with enough inventory to deserve a landing page.
- **Guidance:** questions that require explanation, comparison or expert judgement before a product can be chosen.

Editorial content must earn its existence.

Google's review guidance rewards detailed analysis, evidence of experience, measurements and differentiation for review content. A useful ecommerce editorial programme therefore produces buying evidence - not generic top-of-funnel articles detached from the catalogue. [\[G10\]](#)[\[G11\]](#)

WEAK CONTENT PROGRAMME	STRONGER SEARCH PROGRAMME
"4 blog posts per month"	A backlog of unresolved customer decisions tied to commercial categories.
Generic "best X" pages	First-hand tests, expert comparisons and explicit selection criteria.
Traffic-only reporting	Assisted product discovery, commercial clicks, revenue and repeat search demand.
Content silo	Guides connected to collections/PDPs and commercial pages connected back to guidance.

WARREN'S TAKE

As AI answers absorb more informational clicks, publishing volume for its own sake makes even less sense. Content should teach the customer something that improves the purchase decision, create evidence a search system can reuse, or move authority toward a commercial area. If it does none of those, rankings are not a strategy. [\[W3\]](#)

Authority is product evidence outside your own site.

Digital PR, independent reviews, creator content and specialist references can all shape how a product or brand is perceived during research. The goal is not to seed fake community mentions; it is to create products, evidence, data and stories that credible third parties have a reason to discuss.

PART 11

AI shopping & agentic commerce

Prepare for machine-led discovery without making the future bigger than the present.

AI shopping is becoming a feed problem and an evidence problem.

OpenAI's current Agentic Commerce documentation defines a structured product-feed model with stable product/variant identifiers, titles, descriptions, URLs, brand, price, availability, media and optional policy data. Shopify now exposes eligible merchant products to AI channels through Agentic Storefronts/Shopify Catalog. [\[O1\]\[O2\]\[O3\]\[S6\]\[S7\]](#)

Google is moving in the same direction through Merchant Center. Its `product_detail` and `product_highlight` attributes explicitly state that they can help product discovery across AI-driven surfaces, while Merchant Center's 2026 specification continues adding richer fulfilment and product-data fields. [\[M2\]\[M3\]\[M6\]](#)

What changes for an ecommerce SEO team?

- Product attributes become retrieval dimensions, not just merchandising metadata.
- Variant identity, price and availability need tighter synchronisation across systems.
- Shipping, returns and promotion data matter earlier in discovery because the interface may compare products before a site visit.
- Off-site reviews and brand evidence can influence the research layer even when no classic click occurs.
- Measurement becomes less complete: some discovery, comparison or agent activity can happen before standard analytics sees a session.

WARREN'S TAKE

My priority order is still: win the commercial demand that exists today; make the catalogue and Merchant Center excellent; then extend that product truth into AI/agentic surfaces. The mistake is either extreme - dismissing AI shopping entirely or rebuilding the roadmap around a channel whose demand and attribution are still immature. [\[W2\]\[W4\]](#)

MEASUREMENT WARNING

Agentic shopping can break familiar attribution before it becomes the dominant sales channel. Traffic, CTR and last-click paths may tell only part of the story. Blend platform reporting, product-level visibility, logs, experiments and commercial proxies rather than inventing certainty. [\[W4\]\[SM6\]](#)

THE INDEX CLUB PROCESS

Create an AI-commerce readiness table for the top 100 revenue products: stable ID → product group → attributes → image set → price/availability freshness → shipping/returns → Google feed status → eligible AI channel status → current brand/product representation. Do not scale it catalogue-wide until the model works.

PART 12

International, migrations & scale

Large ecommerce changes should preserve product identity as carefully as URL equity.

Migrations are catalogue transformations, not redirect projects.

A platform migration can change URLs, templates, canonical rules, navigation, variant modelling, product feeds, structured data and international routing in one release. Warren's migration framework emphasises pre-migration benchmarking, full crawls, explicit URL mapping, staging QA, redirect testing and post-launch monitoring. ^[W1]

For ecommerce, map products by product identity where possible.

WARREN'S TAKE

Fuzzy URL matching can accelerate low-risk mapping, but for a large ecommerce migration I would rather match products using a common stable element such as SKU or GTIN when the old and new catalogue share it. URLs can change dramatically while product identity stays exact. ^{[W1][GS1]}

THE INDEX CLUB PROCESS

Build migration joins in layers: **SKU/GTIN exact match → stable product ID → category/taxonomy mapping → URL/path heuristics → manual exception review**. Then validate redirects, selected variants, canonical destinations, feed URLs, structured data and internal links against the final live catalogue.

International stores need explicit locale architecture.

Google recommends locale-specific URLs and hreflang for multi-regional/multilingual sites, warning against relying on IP-based adaptation to expose all versions. Ecommerce adds price, stock, shipping and product-range differences on top, so market templates must be modelled as real product experiences rather than copied translations. ^[G12]

Do not migrate dysfunction perfectly.

The safest migration is not always a one-to-one recreation. A migration is a rare moment when teams can consolidate duplicate collections, clean dead facets, simplify internal linking, repair product identity and remove old platform artefacts. Preserve valuable demand and equity; do not preserve every historical mistake.

PART 13

Measurement & prioritisation

The useful ecommerce SEO metric is the one that changes a commercial decision.

Start with revenue mechanics, not issue counts.

A 900-line technical audit and a five-item revenue backlog are not the same thing. Ecommerce prioritisation works best when the team can connect search opportunity to demand, ranking/visibility gap, expected click capture, conversion rate, average order value or margin, and implementation effort.

A simple opportunity model

ILLUSTRATIVE MODEL

Monthly search demand × expected incremental CTR/click capture × ecommerce CVR × AOV × confidence = directional revenue opportunity. Use ranges, not false precision. Add margin and paid-search dependency when they materially change priority.

This is not a forecast of guaranteed SEO revenue. It is a way to compare opportunities using known commercial inputs rather than ranking every audit item 'high'.

Measure the Catalogue-to-Customer gates separately

GATE	USEFUL SIGNALS
Eligible	Indexation, crawl status, logs, sitemap coverage, feed approvals.
Understood	Identifier completeness, feed/schema consistency, Merchant Center diagnostics.
Discoverable	Non-brand query coverage, category/PDP visibility, internal-search demand coverage.
Competitive	Share of relevant rankings/listings, reviews, authoritative mentions, content evidence.
Selected	Organic clicks, Shopping impressions/free listings, AI citations/recommendations where measurable.
Persuasive	PDP/PLP engagement, add-to-cart, CVR, returns, revenue per landing page.
Measured	Attribution coverage, revenue cohorts, margin, incrementality experiments, paid dependency.

WARREN'S TAKE

I would rather report fewer metrics with an explicit decision attached than give a marketing director another dashboard they cannot act on. Rankings and visibility matter because they are upstream signals. The business outcome is still profitable acquisition and revenue, and AI discovery should not be allowed to turn measurement back into screenshots and vibes.

Use AI visibility carefully.

AI recommendations and product citations are useful directional observations, but prompt-level outputs vary and referral data can be incomplete. The strongest programme uses repeatable prompt panels, first-party platform reporting where available, product-level cohort analysis and experiments. It does not pretend one proprietary score is revenue.

PART 14

The 90-day operating plan

Fix the system in dependency order. Do not start with the most fashionable tactic.

Days 1–30: establish product truth and constraints.

- Baseline organic revenue, non-brand visibility, top commercial landing pages and paid dependency.
- Crawl the site by template; map indexable URL classes, parameters, facets, pagination and internal search.
- Export Merchant Center diagnostics and product data for priority inventory.
- Create the query-to-page map for top categories and products.
- Run feed ↔ page ↔ schema consistency checks on the top 100 products.
- Identify the highest-value category gaps, cannibalisation and broken discovery paths.
- Capture internal-search terms and no-result searches.
- Define the Catalogue-to-Customer failing gate for each major issue.

Days 31–60: fix architecture and money pages.

- Resolve crawl/indexation defects on priority templates.
- Publish/merge/retire categories based on demand + inventory + differentiation + durability.
- Implement the facet and pagination policy.
- Upgrade the highest-opportunity PLPs and PDPs.
- Correct Merchant Center identifiers, variant grouping, prices, availability and rich product attributes.
- Add supported structured data where visible truth and platform eligibility justify it.
- Improve internal linking between home/category/subcategory/PDP and decision content.
- Implement lifecycle rules for OOS and discontinued inventory.

Days 61–90: build evidence, scale and measurement.

- Publish buying/comparison assets around unresolved customer decisions.
- Launch digital PR/authority work tied to commercially important categories.
- Extend structured attributes and feed quality to the wider catalogue.
- Pilot agentic/AI-commerce readiness on priority products where eligible.
- Create a recurring internal-search and feed-diagnostic review.
- Instrument product/category revenue cohorts and relevant AI/shopping signals.
- Build an experiment backlog with a clear expected mechanism and success metric.
- Re-rank the roadmap using impact × confidence × effort and known implementation constraints.

WHAT NOT TO DO FIRST

Do not begin by rewriting every meta description, adding 500 words to every category, generating thousands of “SEO collections”, or buying an AI-visibility tool before the catalogue truth, architecture and commercial priorities are understood.

PART 15

The 150-point ecommerce search audit

Ten systems. Fifteen checks each. Use it as a diagnostic, not a theatre prop.

How to score it.

Score each check **0 = absent/broken**, **1 = partial/inconsistent**, **2 = strong/controlled**. The raw maximum is 300. Do not turn the number into a universal benchmark; use it to expose where a store is weak and where risk is concentrated. Weight commercial areas separately when prioritising action.

SCORING RULE

A technically perfect answer can still be low priority. After scoring, add **Impact / Confidence / Effort** and identify the Catalogue-to-Customer gate each failed check affects.

oi / Crawl & indexation

- | | | | |
|---|---|---|--|
| 0 | 1 | 2 | Priority product, collection and editorial URLs return stable 200 responses. |
| 0 | 1 | 2 | robots.txt does not block pages or resources required for rendering priority templates. |
| 0 | 1 | 2 | No priority templates carry accidental noindex directives. |
| 0 | 1 | 2 | Canonical tags resolve to deliberate, indexable destination URLs. |
| 0 | 1 | 2 | Canonical targets use the preferred protocol, hostname and trailing-slash convention. |
| 0 | 1 | 2 | XML sitemaps contain only canonical, indexable 200 URLs. |
| 0 | 1 | 2 | Product URLs are discoverable through crawlable HTML links, not only internal search or JavaScript events. |
| 0 | 1 | 2 | Pagination exposes crawlable sequential URLs or equivalent crawl paths. |
| 0 | 1 | 2 | Load-more and infinite-scroll implementations have crawlable URL equivalents. |
| 0 | 1 | 2 | Important products are not orphaned when they leave merchandising modules. |
| 0 | 1 | 2 | Parameter and tracking URLs do not create uncontrolled indexable duplicates. |
| 0 | 1 | 2 | Internal search-result URLs have an explicit crawl/indexation policy. |
| 0 | 1 | 2 | Large catalogues have a documented crawl-budget strategy rather than indiscriminate blocking. |
| 0 | 1 | 2 | Server logs or crawl evidence are used to validate bot access for large/complex stores. |
| 0 | 1 | 2 | Indexation is segmented by template and commercial value, not reviewed only as a domain-wide total. |

02 / Architecture & demand

- | | | | |
|---|---|---|---|
| 0 | 1 | 2 | Every major commercial query cluster has a clearly assigned preferred page type. |
| 0 | 1 | 2 | Category hierarchy reflects how customers shop, not only the ERP or internal organisation chart. |
| 0 | 1 | 2 | Broad categories link to durable subcategories before products where the taxonomy warrants it. |
| 0 | 1 | 2 | High-value collection pages are reachable within a sensible number of clicks from strong hubs. |
| 0 | 1 | 2 | No two priority collections exist solely because different teams created them for similar intent. |
| 0 | 1 | 2 | Cannibalisation is checked using query-to-page data, not only keyword-tool overlap. |
| 0 | 1 | 2 | Internal search terms are reviewed for missing categories, attributes and synonyms. |
| 0 | 1 | 2 | Site-search no-result queries are treated as product/discovery research. |
| 0 | 1 | 2 | SERP format is checked before deciding whether a query belongs to a PDP, PLP or guide. |
| 0 | 1 | 2 | Category creation requires meaningful demand, inventory depth and durability. |
| 0 | 1 | 2 | Seasonal pages have a reuse/retention strategy instead of being rebuilt every year. |
| 0 | 1 | 2 | Campaign landing pages do not unnecessarily compete with evergreen category URLs. |
| 0 | 1 | 2 | Brand, product-type and use-case architecture is explicit for multi-brand retailers. |
| 0 | 1 | 2 | Editorial content links into commercial hubs where it genuinely advances the buying task. |
| 0 | 1 | 2 | Architecture decisions are documented so merchandising and development do not undo them later. |

03 / Facets, filters & pagination

- | | | | |
|---|---|---|---|
| 0 | 1 | 2 | Every facet has a documented crawl/indexation decision. |
| 0 | 1 | 2 | Filter combinations cannot generate an effectively infinite crawl space. |
| 0 | 1 | 2 | Sort-order URLs are prevented from creating duplicate indexable pages. |
| 0 | 1 | 2 | Low-value multi-filter combinations remain shopping tools rather than SEO landing pages. |
| 0 | 1 | 2 | High-demand durable facet combinations are promoted into intentional landing pages where warranted. |
| 0 | 1 | 2 | Indexable filtered experiences have self-referencing canonicals and useful unique value. |
| 0 | 1 | 2 | Indexable facet pages are linked internally rather than dependent on crawlable filter links alone. |
| 0 | 1 | 2 | Zero-result filter states cannot become indexable thin pages. |
| 0 | 1 | 2 | Facet values use a controlled taxonomy rather than inconsistent free text. |
| 0 | 1 | 2 | Filters reflect category-specific decision attributes, not only generic price/brand/colour. |
| 0 | 1 | 2 | Applied filters are visible and removable for users. |
| 0 | 1 | 2 | Multi-select behaviour matches the way customers actually refine the category. |
| 0 | 1 | 2 | Pagination URLs are unique, crawlable and internally linked. |
| 0 | 1 | 2 | Each paginated URL contains the correct products without requiring a user action to reveal them. |
| 0 | 1 | 2 | Products beyond page one remain discoverable without relying solely on sitemaps. |

04 / Category & PLP quality

- | | | | |
|---|---|---|---|
| 0 | 1 | 2 | Each priority PLP has a clear H1 aligned with the product set and search intent. |
| 0 | 1 | 2 | Title tags differentiate category scope and avoid boilerplate duplication. |
| 0 | 1 | 2 | Intro copy explains the collection rather than padding the page with generic SEO prose. |
| 0 | 1 | 2 | Critical buying information is visible before excessive editorial copy pushes products down. |
| 0 | 1 | 2 | Product cards expose the attributes customers need to compare options. |
| 0 | 1 | 2 | Pricing and availability are accurate on PLP cards where displayed. |
| 0 | 1 | 2 | Filters use terminology customers understand. |
| 0 | 1 | 2 | Category-specific filters exist for the important decision attributes. |
| 0 | 1 | 2 | Products are merchandised so strong/representative inventory appears early without hiding relevance. |
| 0 | 1 | 2 | Related subcategories and adjacent collections are linked contextually. |
| 0 | 1 | 2 | Breadcrumbs reflect the intended hierarchy and are crawlable. |
| 0 | 1 | 2 | Empty categories have a deliberate UX and indexation response. |
| 0 | 1 | 2 | Low-stock categories are monitored before they become poor search destinations. |
| 0 | 1 | 2 | Helpful buying guidance is available where the decision genuinely needs explanation. |
| 0 | 1 | 2 | PLP performance is measured against both organic acquisition and product-finding/conversion outcomes. |

05 / Product detail pages

0	1	2	Product names use shopper language and preserve the real product identity.
0	1	2	PDP copy explains what the product is, who it is for and the important constraints.
0	1	2	Important specifications are present as text/data, not trapped only in images or PDFs.
0	1	2	Variant options expose correct price, availability, imagery and identifiers.
0	1	2	Each important variant has a stable deep-linkable state where the platform supports it.
0	1	2	Product images are high quality, crawlable and supported by descriptive alt text.
0	1	2	PDPs include useful visual evidence such as additional angles, scale, fit or use context where relevant.
0	1	2	Shipping expectations are clear before checkout.
0	1	2	Return policy information is easy to find and internally consistent.
0	1	2	Reviews provide genuine decision support rather than decorative star counts.
0	1	2	FAQs answer product-specific objections rather than repeating generic company information.
0	1	2	Related products and accessories are contextually useful and crawlable.
0	1	2	Comparison/help content is linked where customers commonly hesitate between options.
0	1	2	Out-of-stock messaging provides useful next actions rather than a dead end.
0	1	2	PDP success is measured in revenue/conversion as well as impressions and rankings.

o6 / Product identity, feeds & structured data

☐ ☐ ☐ Merchant Center is connected and treated as organic-discovery infrastructure, not only a paid-media feed.

☐ ☐ ☐ Product feed IDs are stable and never casually recycled for different products.

☐ ☐ ☐ GTINs are supplied when valid manufacturer identifiers exist.

☐ ☐ ☐ GTINs are never invented to fill a field.

☐ ☐ ☐ Brand values reflect the actual manufacturer/brand rather than the retailer where inappropriate.

☐ ☐ ☐ Variant items share a stable group identifier.

☐ ☐ ☐ Feed title, landing-page title and selected variant describe the same purchasable item.

☐ ☐ ☐ Feed price matches the landing-page price for the selected offer.

☐ ☐ ☐ Feed availability matches the landing page and backend stock state.

☐ ☐ ☐ Product schema reflects visible on-page truth.

☐ ☐ ☐ ProductGroup/variant markup matches the actual variant model.

☐ ☐ ☐ Shipping and returns data is supplied in the appropriate feed/schema/account location.

☐ ☐ ☐ Product highlights describe useful benefits rather than keyword stuffing.

☐ ☐ ☐ Product detail fields capture technical, verifiable attributes where useful.

☐ ☐ ☐ Feed diagnostics are part of SEO QA and are not left solely to paid media.

07 / Inventory & product lifecycle

- | | | | |
|---|---|---|--|
| 0 | 1 | 2 | New products have a planned discovery path before launch. |
| 0 | 1 | 2 | New PDPs are linked from relevant collections immediately when they become commercially useful. |
| 0 | 1 | 2 | Temporarily out-of-stock products remain live when the product is expected to return and the page still serves demand. |
| 0 | 1 | 2 | Temporary OOS pages clearly communicate status and offer relevant alternatives or restock actions. |
| 0 | 1 | 2 | Permanently discontinued products are distinguished from temporarily unavailable products. |
| 0 | 1 | 2 | Discontinued products with a true successor have a deliberate redirect/replacement strategy. |
| 0 | 1 | 2 | Discontinued products without an equivalent are not blindly redirected to unrelated categories. |
| 0 | 1 | 2 | Product removal rules consider traffic, backlinks, demand, replacement and customer utility. |
| 0 | 1 | 2 | Seasonal products have documented off-season handling. |
| 0 | 1 | 2 | Returning seasonal PDPs reuse durable URLs where appropriate. |
| 0 | 1 | 2 | Variant discontinuation does not break the parent product experience. |
| 0 | 1 | 2 | Feeds remove or update discontinued inventory consistently with the website. |
| 0 | 1 | 2 | Sitemaps stop advertising dead inventory as current. |
| 0 | 1 | 2 | Internal links are updated when a product leaves the catalogue permanently. |
| 0 | 1 | 2 | Lifecycle decisions preserve valuable information when the page still satisfies research demand. |

o8 / Content, expertise & internal search

- | | | | |
|---|---|---|--|
| 0 | 1 | 2 | Content strategy begins with customer decisions, not monthly blog quotas. |
| 0 | 1 | 2 | Buying guides resolve real product-selection questions and link to relevant commercial pages. |
| 0 | 1 | 2 | Comparison content demonstrates criteria and evidence rather than generic summaries. |
| 0 | 1 | 2 | Editorial recommendations disclose relevant commercial relationships. |
| 0 | 1 | 2 | First-hand tests include original observations, images or measurements where possible. |
| 0 | 1 | 2 | Experts or qualified staff are attributed where judgement is central to the recommendation. |
| 0 | 1 | 2 | Internal search vocabulary informs category, synonym and content decisions. |
| 0 | 1 | 2 | On-site search tolerates common misspellings and customer language where the platform allows. |
| 0 | 1 | 2 | No-result searches feed a recurring merchandising/content backlog. |
| 0 | 1 | 2 | Blog/editorial pages do not cannibalise commercial pages by targeting the same task with the wrong format. |
| 0 | 1 | 2 | High-performing informational pages contain deliberate routes to relevant product discovery. |
| 0 | 1 | 2 | Commercial pages link back to deeper guidance where it reduces uncertainty. |
| 0 | 1 | 2 | Content decay is monitored for product availability, outdated models and superseded recommendations. |
| 0 | 1 | 2 | AI-generated copy is human-reviewed for product truth, differentiation and legal accuracy. |
| 0 | 1 | 2 | Automation is used for consistency and scale; important pages receive human judgement. |

09 / Authority, reviews & corroboration

0	1	2	Priority commercial pages earn links/mentions through genuinely useful assets, products or stories.
0	1	2	Digital PR is connected to categories/products that have commercial upside, not treated as a disconnected vanity channel.
0	1	2	Brand mentions are monitored alongside backlinks.
0	1	2	Review acquisition follows platform policies and never relies on fabrication or incentives that breach rules.
0	1	2	Review content is visible and useful to customers.
0	1	2	Review markup matches reviews actually shown on the page.
0	1	2	Product review editorial demonstrates first-hand or expert evidence where claimed.
0	1	2	Media coverage reinforces real product claims rather than invented superlatives.
0	1	2	Creator/influencer content is considered part of the wider product evidence ecosystem.
0	1	2	Manufacturer and retailer product information are consistent where they refer to the same item.
0	1	2	Priority brands/products have credible independent references beyond the merchant site.
0	1	2	PR/product launches preserve stable URLs and reusable evidence after the campaign ends.
0	1	2	Comparison/review queries are monitored for inaccurate third-party narratives.
0	1	2	Leadership/expert commentary is used where it creates differentiated category expertise.
0	1	2	Authority work is measured against commercial visibility and assisted outcomes, not link count alone.

io / AI, agentic commerce & measurement

0	1	2	AI shopping is treated as an emerging discovery layer, not a replacement for Google/Search/Shopping fundamentals.
0	1	2	Google Merchant Center product data is complete enough to support richer and AI-driven discovery surfaces.
0	1	2	OpenAI/Shopify agentic-channel eligibility is understood where commercially relevant.
0	1	2	OpenAI product-feed requirements are mapped against the store's current catalogue model where applicable.
0	1	2	Product identifiers remain stable across website, feed and agentic integrations.
0	1	2	Prices and availability update frequently enough for the volatility of the catalogue.
0	1	2	Shipping/returns/product-policy data is machine-readable and human-readable.
0	1	2	AI/search monitoring separates brand mention, product citation, recommendation and referral outcomes.
0	1	2	Agent-driven journeys are not assumed to be fully visible in standard web analytics.
0	1	2	Log-level and platform reporting are considered where first-party referral data is incomplete.
0	1	2	Organic revenue is reported by landing-page/product/category cohort where possible.
0	1	2	SEO opportunity models incorporate current demand, conversion rate, AOV/margin and paid dependency.
0	1	2	Paid and organic overlap is reviewed to identify expensive demand the brand could earn more efficiently.
0	1	2	Experiments define the expected commercial mechanism before implementation.
0	1	2	The roadmap is prioritised by impact, confidence and effort rather than by audit-error count.

PART 16

Myth book

Popular ecommerce SEO advice, with the missing context put back in.

01

Every product description must be unique.

Unique wording is not the same as unique value. Manufacturer copy can be insufficient, but the real goal is useful product truth, decision attributes and evidence. Add what the retailer uniquely knows instead of forcing cosmetic rewrites across thousands of low-value SKUs.

02

Out-of-stock products should always 301.

Temporary OOS, permanent discontinuation and a true replacement are different states. Redirect only when another URL genuinely satisfies the same intent.

03

All faceted URLs are bad for SEO.

Most facet combinations should stay shopping controls, but some durable combinations map to real demand. Promote those deliberately; do not confuse blanket blocking with strategy.

04

More indexed products means more organic growth.

Indexation is eligibility, not demand. Large stores can add thousands of weak or duplicate URLs without creating more useful search coverage.

05

Category pages need 500 words of copy.

There is no magic word count. A category needs enough information to orient the shopper, explain the set and support the decision without burying the products.

06

Schema makes products rank higher.

Structured data can improve understanding and eligibility for supported search features. It is not a generic ranking-boost switch. [\[G4\]](#)[\[G5\]](#)

07

The blog is the ecommerce SEO strategy.

Editorial content is useful when it resolves buying questions, earns authority or strengthens commercial discovery. Commercial architecture and product data usually deserve attention first.

08

Internal search pages should rank because customers search those terms.

Internal-search data is excellent demand research. The raw result URL is often a poor permanent landing page; durable demand should usually graduate into an intentional category or guide.

09

AI shopping means product pages matter less.

AI/agent systems still need accurate product information, URLs, offers and evidence. The PDP remains a critical source of truth and conversion interface even when discovery happens elsewhere.

IO

Feed work belongs to paid media.

Merchant Center affects free/product surfaces and product understanding. Feed defects can be organic discovery defects. [\[G3\]](#) [\[M1\]](#)

II

A migration is mostly a redirect map.

For ecommerce it can also change product identity, variants, feeds, canonicals, navigation and market structure. Redirects are one dependency inside a larger catalogue transformation.

I2

Automation can replace ecommerce SEO judgement.

Automation is excellent for repetitive mapping, classification and QA. It is weakest at deciding whether a page should exist, which intent it owns, and whether the commercial trade-off is worthwhile. [\[W5\]](#)

PART 17

Source library

Credited evidence, standards, practitioner research and the Warren/Delphi layer.

67 sources reviewed and credited.

This library is intentionally mixed. Platform documentation is used for supported behaviour and eligibility. Standards define product identity and vocabulary. UX research informs product-finding and decision design. Practitioner sources help with implementation patterns. Warren/Delphi material is used as a judgement layer - not as external proof of how a platform works.

URLs were current at the research snapshot. Platform documentation changes; re-check implementation requirements before shipping production changes.

G1	PLATFORM SEO best practices for ecommerce sites Google Search Central https://developers.google.com/search/docs/specialty/ecommerce
G2	PLATFORM Help Google understand your ecommerce site structure Google Search Central https://developers.google.com/search/docs/specialty/ecommerce/help-google-understand-your-ecommerce-site-structure
G3	PLATFORM Share your product data with Google Google Search Central https://developers.google.com/search/docs/specialty/ecommerce/share-your-product-data-with-google
G4	PLATFORM Merchant listing structured data Google Search Central https://developers.google.com/search/docs/appearance/structured-data/merchant-listing
G5	PLATFORM Product snippet structured data Google Search Central https://developers.google.com/search/docs/appearance/structured-data/product-snippet
G6	PLATFORM Product variant structured data Google Search Central https://developers.google.com/search/docs/appearance/structured-data/product-variants
G7	PLATFORM Pagination, incremental page loading and Google Search Google Search Central https://developers.google.com/search/docs/specialty/ecommerce/pagination-and-incremental-page-loading
G8	PLATFORM Designing a URL structure for ecommerce sites Google Search Central https://developers.google.com/search/docs/specialty/ecommerce/designing-a-url-structure-for-ecommerce-sites
G9	PLATFORM Google Images SEO best practices Google Search Central https://developers.google.com/search/docs/appearance/google-images

G10

PLATFORM

Google Search reviews system

Google Search Central

<https://developers.google.com/search/docs/appearance/reviews-system>

G11

PLATFORM

Write high quality reviews

Google Search Central

<https://developers.google.com/search/docs/specialty/ecommerce/write-high-quality-reviews>

G12

PLATFORM

Managing multi-regional and multilingual sites

Google Search Central

<https://developers.google.com/search/docs/specialty/international/managing-multi-regional-sites>

G13

PLATFORM

Link best practices for Google

Google Search Central

<https://developers.google.com/search/docs/crawling-indexing/links-crawlable>

G14

PLATFORM

Optimize your crawl budget

Google Crawling Infrastructure

<https://developers.google.com/crawling/docs/crawl-budget>

G15

PLATFORM

Google Search Essentials

Google Search Central

<https://developers.google.com/search/docs/essentials>

G16

PLATFORM

Where ecommerce data can appear on Google

Google Search Central

<https://developers.google.com/search/docs/specialty/ecommerce/where-ecommerce-data-can-appear-on-google>

M1

PLATFORM

Create a product file for Merchant Center

Google Merchant Center

<https://support.google.com/merchants/answer/12631822?hl=en>

M2

PLATFORM

Product detail [product_detail]

Google Merchant Center

<https://support.google.com/merchants/answer/9218260?hl=en>

M3

PLATFORM

Product highlight [product_highlight]

Google Merchant Center

<https://support.google.com/merchants/answer/9216100?hl=en>

M4

PLATFORM

Supported structured data attributes and values

Google Merchant Center

<https://support.google.com/merchants/answer/6386198?hl=en>

M5

PLATFORM

GTIN [gtin]

Google Merchant Center

<https://support.google.com/merchants/answer/6324461?hl=en>

M6

PLATFORM

Merchant Center product data specification update 2026

Google Merchant Center

<https://support.google.com/merchants/answer/16989427?hl=en-GB>

M7

PLATFORM

About sale price annotations

Google Merchant Center

<https://support.google.com/merchants/answer/9017019?hl=en>

M8

PLATFORM

Sale price [sale_price]

Google Merchant Center

<https://support.google.com/merchants/answer/6324471?hl=en>

M9

PLATFORM

Make your products stand out and maximise reach

Google Merchant Center

<https://support.google.com/merchants/answer/6167176?hl=en>

O1

PLATFORM

Agentic Commerce: product feed specification

OpenAI Developers

<https://developers.openai.com/commerce/specs/file-upload/products>

O2

PLATFORM

Agentic Commerce API overview

OpenAI Developers

<https://developers.openai.com/commerce/specs/api/overview>

O3

PLATFORM

Agentic Commerce API: products

OpenAI Developers

<https://developers.openai.com/commerce/specs/api/products>

O4	PLATFORM Agentic Commerce API: promotions OpenAI Developers https://developers.openai.com/commerce/specs/api/promotions
O5	PLATFORM Merchant Feed Terms of Service OpenAI https://openai.com/policies/merchant-feed-terms-of-service/
O6	PLATFORM Shopping from Shopify merchants in ChatGPT OpenAI Help https://help.openai.com/en/articles/20001241-shopping-from-shopify-merchants-in-chatgpt
S1	PRACTITIONER / PLATFORM SEO basics for ecommerce (2026) Shopify https://www.shopify.com/blog/seo-basics
S2	PRACTITIONER / PLATFORM Ecommerce SEO audit: Shopify checklist (2026) Shopify https://www.shopify.com/blog/ecommerce-seo-audit
S3	PRACTITIONER / PLATFORM Faceted navigation: SEO and UX best practices Shopify https://www.shopify.com/blog/faceted-navigation
S4	PRACTITIONER / PLATFORM Internal linking benefits (2026) Shopify https://www.shopify.com/blog/internal-linking-benefits
S5	PRACTITIONER / PLATFORM Technical SEO for ecommerce Shopify https://www.shopify.com/blog/technical-seo
S6	PLATFORM Shopify agentic storefronts Shopify Help https://help.shopify.com/en/manual/online-sales-channels/agentic-storefronts

S7	PLATFORM Selling on ChatGPT Shopify Help https://help.shopify.com/en/manual/online-sales-channels/agentive-storefronts/chatgpt
H1	PRACTITIONER Ecommerce SEO: a beginner's guide Ahrefs https://ahrefs.com/blog/ecommerce-seo/
H2	PRACTITIONER Ways to improve ecommerce category pages for SEO Ahrefs https://ahrefs.com/blog/seo-ecommerce-category-pages/
SM1	PRACTITIONER SEO ecommerce category pages: optimisation guide Semrush https://www.semrush.com/blog/seo-ecommerce-category-pages/
SM2	PRACTITIONER Ecommerce marketing: strategies for search and AI in 2026 Semrush https://www.semrush.com/blog/ecommerce-marketing/

SM3

PRACTITIONER

Ecommerce AI SEO: optimise online stores for LLMs

Semrush

<https://www.semrush.com/blog/ecommerce-ai-seo/>

SM4

PRACTITIONER

Complete ecommerce SEO guide

Semrush

<https://www.semrush.com/blog/ecommerce-seo/>

SM5

PRACTITIONER

Deep dive into SEO for ecommerce product pages

Semrush

<https://www.semrush.com/blog/ecommerce-product-page-seo/>

SM6

INDUSTRY REPORTING

Google to add AI performance report to Merchant Center

Semrush

<https://www.semrush.com/blog/google-to-add-ai-performance-report-to-merchant-center/>

B1

UX RESEARCH

Product listing page UX: best practices (2026)

Baymard Institute

<https://baymard.com/blog/product-listing-page-plp-ux>

B2

UX RESEARCH

Ecommerce filters: UI best practices (2026)

Baymard Institute

<https://baymard.com/blog/ecommerce-filter-ui>

B3

UX RESEARCH

Product lists and filtering UX research

Baymard Institute

<https://baymard.com/research/ecommerce-product-lists>

B4

UX RESEARCH

Ecommerce category pages: design tips (2026)

Baymard Institute

<https://baymard.com/blog/ecommerce-category-page>

B5

UX RESEARCH

Ecommerce UX audit: product discovery

Baymard Institute

<https://baymard.com/blog/ecommerce-ux-audit>

B6

UX RESEARCH

Product page UX: pitfalls and best practices (2026)

Baymard Institute

<https://baymard.com/research-articles/current-state-ecommerce-product-page-ux>

B7

UX RESEARCH

Findability and discoverability for ecommerce

Baymard Institute

<https://baymard.com/blog/findability-vs-discoverability-ux>

B8

UX RESEARCH

Top ecommerce sites UX benchmark

Baymard Institute

<https://baymard.com/ux-benchmark>

W0

PRACTITIONER CONTEXT

Delphi - Warren Hance

Warren Hance

<https://www.delphi.ai/warren-hance>

W1

PRACTITIONER

The ultimate SEO guide to website migrations and mergers

Warren Hance

<https://warrenhanceseo.com/the-ultimate-seo-guide-to-website-migrations-and-mergers/>

W2

PRACTITIONER

Comprehensive playbook for appearing in ChatGPT Shopping (2026)

Warren Hance

<https://warrenhanceseo.com/comprehensive-playbook-for-appearing-in-chatgpt-shopping-2026/>

W3

PRACTITIONER

AI, ecommerce and SEO: hybrid discovery and conversion

Warren Hance

https://www.linkedin.com/posts/warren-hance-86869b122_last-week-in-ai-ecommerce-seo-edition-activity-7393500782098223104-tAwA

W4

PRACTITIONER

2026 ecommerce SEO prediction: agentic shopping and measurement

Warren Hance

<https://productrise.app/blog/e-commerce-seo-predictions-2026>

W5

PRACTITIONER

Automation with a human judgement layer

Warren Hance

https://www.linkedin.com/posts/warren-hance-86869b122_most-people-build-h1s-manually-id-rather-activity-7399233814780133376-w6qG

W6

PRACTITIONER

SEO tools, scripts and practical resources

Warren Hance

<https://warrenhanceseo.com/seo-tools-scripts-practical-resources/>

SC1

STANDARD

Product

Schema.org

<https://schema.org/Product>

SC2

STANDARD

ProductGroup

Schema.org

<https://schema.org/ProductGroup>

SC3

STANDARD

Offer

Schema.org

<https://schema.org/Offer>

SC4

STANDARD

MerchantReturnPolicy

Schema.org

<https://schema.org/MerchantReturnPolicy>

GS1

STANDARD

Global Trade Item Number (GTIN)

GS1

<https://www.gs1.org/standards/id-keys/gtin>

GS2

STANDARD

GS1 in retail: product data journey

GS1

<https://www.gs1.org/industries/retail>

FIELD MANUAL 02

Make the catalogue earn the customer.

Search visibility is not a metadata project. It is the cumulative result of product truth, architecture, evidence, discovery and a buying experience that holds up after the click.

[THEINDEXCLUB.CO.UK](https://theindexclub.co.uk)