

Example Shop

Can Google, AI answer engines, and assistive technology
find, understand, and recommend your store?

Client: Example Shop · **Domain:** example-shop.com

Platform detected: Shopify · **Market:** United States (primary) + EU / UK / CA

Date: 18 July 2026

Author: Yassin BELHAJ SALAH

Yassin BELHAJ SALAH, E-commerce Solutions Architect

Confidentiality: Internal use — Example Shop

© Yassin BELHAJ SALAH — 2026. All rights reserved.


Executive summary

Scope & method. This is a **preliminary external evaluation** — it relies only on publicly observable signals (rendered HTML, HTTP headers, sitemaps, Lighthouse / PageSpeed Insights), **without any internal access** to your code, admin, analytics or Search Console. It is a diagnostic assisted by an AI workflow, with manual validation of the retained findings. Every conclusion is scoped to the **pages actually tested**: the homepage, one collection page (/collections/earbuds) and one product page (/products/crusher-1080-anc-headphones). Where a conclusion would require internal access, it is written as a hypothesis to validate — never as certainty.

Headline: you have an excellent discoverability foundation — this is a refinement play, not a rebuild. Your product pages already expose a valid Offer (price, currency, in-stock), a rich ProductGroup with variants, FAQ structured data, a full 25-locale hreflang cluster, clean self-referencing canonicals, and even the brand-new Shopify agentic-commerce layer. Accessibility sits at market parity. The value on the table is in a handful of precise details that convert this foundation into **rich results** in Google and **citations** in AI answers.

Discoverability radar

Scores /10 derived from the measured signals below — not opinion. A low score maps to a concrete, verifiable observation.



ACCESSIBILITY (LIGHTHOUSE / PSI MOBILE)

93/100

HREFLANG LOCALES (INCL. X-DEFAULT)

25

PDP AGGREGATERATING VALUE

null

The five things that matter most

#	What we observed (tested pages)	Why it matters	Priority
1	PDP emits an AggregateRating with ratingValue: null & reviewCount: null, while review widgets (Loox / Okendo) render on the page.	Invalid rating object → risk of losing review-star rich results and of "invalid" warnings in Search Console.	Critical
2	No BreadcrumbList structured data on the tested PDP or collection page.	Loses breadcrumb rich results and weakens hierarchy / entity signals used by AI answer engines.	High
3	Accessibility 93/100, with color-contrast (×3) and link-name (×3) failing; manual WCAG layer untested.	Good machine layer, but the EU market you serve is covered by the European Accessibility Act.	High
4	PDP <title> is 155 characters; multiple <h1> per page (home 3, PDP 2).	SERP truncation and diluted semantic hierarchy for crawlers and LLMs.	Medium
5	The Shopify agentic files (agents.md, llms.txt, .well-known/ucp) are live but are the generic platform default.	An opportunity (prospective): customize them so AI agents get your brand, not boilerplate.	Prospective

Prioritisation is a strategic judgement, not a measurement. Prospective items (GEO / agentic files) sit deliberately behind the measurable structured-data and accessibility work.

What already works — and the platform underneath

A credible plan starts by recognising the solid base you have built. On the tested pages, the following are genuinely well done:

Valid product Offer. The PDP Offer carries availability: InStock, price: 279.99, priceCurrency: USD and a canonical URL — exactly what Google needs for product results.

Real international setup. A 25-locale hreflang cluster with x-default, powered by Shopify Markets (not a Google Translate widget). This is an asset most competitors do not match.

Rich page types. Product + ProductGroup (with colour variants), FAQPage on both PDP and collection, Organization + WebSite on the homepage, complete Open Graph.

Clean crawl surface. A sitemap index with dedicated product, collection, page and blog sitemaps; self-referencing canonicals on every tested template; an accessibility score at market parity (93/100).

Native Shopify discoverability features — audit

Method: capabilities confirmed against what your storefront actually returns today, cross-checked with current Shopify documentation. We recommend only native / free levers — no paid app.

Native feature	Status (observed)	Leverage
Automatic XML sitemap (index + products/collections/pages/blogs)	Active	Confirm every child sitemap is submitted in Search Console & Bing Webmaster Tools.
Self-referencing canonical tags	Active	Confirmed on home, PDP and collection — keeps duplicate/variant URLs from diluting equity.
hreflang via Shopify Markets	Active	25 locales + x-default. Add <code>og:locale:alternate</code> for cleaner social sharing per region.
Theme JSON-LD (Product / Offer / FAQ / Organization / WebSite)	Partial	Offer is valid, but the <code>AggregateRating</code> is null and <code>Breadcrumb / ItemList</code> are absent — the core refinement (CT-01, CT-02, CT-08).
Search engine listing (title / meta templates)	Active	Descriptive titles & metas in place; trim the over-long PDP title (CT-04).
Customisable <code>robots.txt.liquid</code>	Active	Standard Shopify rules + Ahrefs/MJ12/Nutch blocked. Add a documented AI-crawler policy (CT-10).
Agentic layer (<code>llms.txt</code> / <code>agents.md</code> / <code>.well-known/ucp</code> , UCP 2026-04-08)	Active generic	Shopify default — personalise via Liquid so it carries your brand (CT-09, prospective).
Per-image alt-text fields (admin)	Partial	19/42 homepage images carry an <code>alt</code> attribute — classify decorative vs informative, then complete (CT-11).

Third-party review apps (Loox, Okendo on the PDP; a Yotpo signal on the collection page) are detected in the markup — active usage and configuration not confirmed from the outside. Your Shopify plan/edition is not observable externally and is not assumed here.

SEO — on-page foundations

On-page signals tell Google (and every crawler) what each page is about. Yours are largely healthy; the fixes below are template-level and fast.

Signal	Home	Collection (earbuds)	PDP (Crusher 1080 ANC)
Descriptive <title>	Yes	Yes	Yes but 155 chars
Meta description	Present	Present	Present
Single semantic <h1>	3 h1	1 h1	2 h1
Self-referencing canonical	Yes	Yes	Yes
Open Graph complete	Yes	Yes	Yes (+price)
Twitter/X Card image	Missing	—	—

CT-04 : Trim the over-long product title

Observation. The tested PDP <title> is 155 characters ("Crusher 1080 ANC - Bass Headphones + Sound by Bose + Bose Quiet Control Noise Canceling | Wireless | Bluetooth 5.3 | Bui..."). **Risk.** Google typically renders ~55-60 characters; the rest is truncated, so the differentiating benefit is often cut off in the SERP. **Solution.** Adjust the product title template to front-load the model + one key benefit, keeping the brand suffix short. Validate the rendered length on a sample of PDPs.

Profile: SEO | **Seniority:** Confirmed | **Effort:** 1 man-day

Expected gain: a fully visible, benefit-led SERP snippet on every product page — higher click-through on the same ranking position.

CT-05 : One semantic <h1> per page

Observation. The homepage renders 3 <h1> ("FEEL SOMETHING REAL", "INTRODUCING CRUSHER 1080 ANC", ...) and the tested PDP renders 2 identical <h1> ("Crusher® 1080 ANC"). **Risk.** Multiple top-level headings dilute the "what is this page about" signal for crawlers and, increasingly, for LLMs that lean on heading structure to extract answers. **Solution.** Keep one <h1> per page (the page's true subject); demote the others to <h2>/<h3> in the theme sections.

Profile: Frontend | **Seniority:** Confirmed | **Effort:** 1 man-day

Expected gain: a cleaner heading tree — a shared benefit for SEO, accessibility (screen-reader navigation) and AI extraction.

CT-07 : Complete the social-card metadata

Observation. The homepage exposes a Twitter/X Card (card, title, description, site) but **no** `twitter:image`, and no `og:locale:alternate` despite 25 hreflang locales. **Risk.** Minor — X usually falls back to `og:image` — but explicit tags avoid inconsistent previews and help per-region sharing. **Solution.** Add `twitter:image` (reuse the OG image) and `og:locale:alternate` for the active locales in the theme <head>.

Profile: Frontend | **Seniority:** Junior | **Effort:** 0.5 man-day

Expected gain: consistent, branded link previews across social and messaging platforms.

Structured data — the highest-value fixes

Structured data (Schema.org / JSON-LD) is the machine-readable layer Google uses for rich results and that AI answer engines lean on to understand entities. Your PDP is *rich* — which makes the two gaps below unusually high-leverage, because everything around them is already right.

CT-01 — Product AggregateRating is emitted with null values. On the tested PDP (/products/crusher-1080-anc-headphones), the server-rendered Product and ProductGroup JSON-LD each contain:

```
"aggregateRating": {
  "@type": "AggregateRating",
  "ratingValue": null,      // no value
  "reviewCount": null      // no value
}
```

...yet review widgets (Loox / Okendo) render on the page and the star rating is drawn client-side.

Risk. An AggregateRating without ratingValue is invalid: it can trigger "invalid object" notices in Search Console and, more importantly, makes the page ineligible for the review-star rich result that other audio brands can earn. AI engines that read the server HTML also receive no rating signal.

Solution. Wire the review app's aggregate rating into the theme's Product schema (Loox and Okendo both expose the aggregate via metafields / Liquid), or remove the empty node and let the app inject a single valid one. **Then validate the rendered page** with Google's Rich Results Test and Search Console URL Inspection — do not assume the theme change alone is sufficient.

Profile: Dev SEO | **Seniority:** Senior | **Effort:** 1.5 man-days

Expected gain: eligibility for review stars in the SERP. Quantify with your data — $Gain \approx (uplift\ in\ click-through\ from\ stars) \times (product-page\ impressions) \times conversion \times average\ order\ value$ — a 30-minute scoping workshop with your Search Console + Analytics turns this into a number.

CT-02 : Add BreadcrumbList structured data

Observation. Neither the tested PDP nor the collection page emits BreadcrumbList / itemListElement (only a CSS "breadcrumb" class is present). **Risk.** No breadcrumb rich result, and a weaker site-hierarchy signal for both Google and AI answer engines that use it to place a product within a catalogue. **Solution.** Add a BreadcrumbList JSON-LD block (Home → Collection → Product) via the theme; validate with the Rich Results Test.

Profile: Dev SEO | **Seniority:** Confirmed | **Effort:** 1 man-day

Expected gain: breadcrumb display in the SERP and clearer catalogue structure for machines.

CT-06 : Enrich the Organization entity

Observation. The homepage Organization has a sameAs array (social profiles) but includes a few empty "" entries, and **no ContactPoint**. For reference, Beats exposes a ContactPoint in its Organization schema. Note also that Google's SearchAction sitelinks search box was **retired on 2024-11-21** — your WebSite/SearchAction markup remains useful semantic metadata for entity understanding, but it no longer produces that old SERP feature. **Solution.** Add a ContactPoint (customer-service phone / URL / languages) and clean the empty sameAs entries — cheap entity-clarity wins that help the Knowledge Graph and AI engines identify your brand unambiguously.

Profile: SEO | **Seniority:** Junior | **Effort:** 0.5 man-day

Expected gain: stronger, unambiguous brand-entity recognition across Google and AI answers.

CT-08 : Structure the collection pages

Observation. The tested collection page carries FAQPage + Organization but no product ItemList / CollectionPage. **Risk.** Category pages are strong ranking assets; without an ItemList, their product grid is less machine-legible. **Solution.** Add CollectionPage + ItemList JSON-LD summarising the listed products via the theme.

Profile: Dev SEO | **Seniority:** Confirmed | **Effort:** 1 man-day

Expected gain: richer, more legible category pages for search and AI.

GEO — visibility in AI answer engines (prospective)

A growing share of product research now happens inside ChatGPT, Perplexity, Google AI Overviews, Gemini and Claude, which cite a few trusted sources instead of listing ten links. Generative Engine Optimization (GEO) is about being one of those cited sources. Per 2026 GEO analyses (LLMrefs' *2026 GEO Guide*; Omnibound's *2026 GEO Statistics*), the levers are machine readability, entity clarity, valid structured data and clear Q&A structure — content with rich structured data and FAQ/Q&A tends to appear more often in AI answers (an industry-level, directional signal, not a measurement of your site).

Honest framing. The single biggest GEO input you control is the same **measurable** structured data covered above (a valid `AggregateRating`, `BreadcrumbList`, `FAQPage`, a clean `Organization`). Rich, valid structured data *improves machine readability and may strengthen your eligibility and representation* — it does **not** guarantee that an AI engine will cite the page. That is why GEO is ranked deliberately *behind* the structured-data work, not ahead of it. Hierarchy here is strategic judgement, not a measurement.

You are unusually well-positioned — because of Shopify

All four agentic-discovery endpoints are live on your store and correctly typed:

- `/llms.txt` & `/llms-full.txt` — **200** text/markdown
- `/agents.md` — **200** text/markdown (Shopify's canonical agent file)
- `/.well-known/ucp` — **200** JSON, Universal Commerce Protocol v2026-04-08, with an MCP endpoint and a checkout capability
- `/sitemap_agentic_discovery.xml` — **200**

This is a real, current asset — but an honest one: Shopify **silently rolled these out across all active stores in May 2026** and turned UCP on by default on **17 June 2026** (Shopify Engineering, *Building the UCP*, 2026). Your files today are the **generic platform default**, personalised only with your store name; the body points agents to the shared `shop.app` skill rather than to your brand story or curated catalogue. Crucially, Shopify makes these files **customisable via Liquid theme templates**.

CT-09 : Customise the agentic files (prospective)

Solution. Override `agents.md` / `llms.txt` in the theme to carry a concise brand narrative ("We Are Example Shop..."), your hero product lines (Crusher, Dime, Sesh, gaming), category entry points and shipping/returns basics — so an AI agent represents you accurately rather than serving boilerplate. Keep valid PDP structured data as the real citation input.

Profile: **SEO** | **Content** | **Seniority:** **Confirmed** | **Effort:** 2 man-days · *prospective, after the measurable work*


CT-10 : Decide and document an AI-crawler policy (prospective)

Observation. Your `robots.txt` (167 lines) blocks SEO crawlers (Ahrefs, MJ12, Nutch) but contains **no explicit directive for AI crawlers** (GPTBot, ClaudeBot, PerplexityBot, Google-Extended, CCBot) — they are governed by the wildcard rules. **Solution.** Make this a deliberate business decision (visibility in AI answers vs. control of content use), then document it explicitly in `robots.txt.liquid`. There is no single "right" answer — it depends on your brand strategy.

Profile: **SEO** | **Seniority:** **Confirmed** | **Effort:** 0.5 man-day · *prospective*

Accessibility (a11y) — measured, plus the legal frame

Accessibility and SEO read the same raw material: the semantics of the HTML. What a screen reader can parse, a crawler and an LLM can parse too — so this section is a double lever.



Read the score honestly. An automated audit like Lighthouse covers only ~30-40% of WCAG. 93/100 is genuinely good *for the machine-testable subset*, but keyboard operability, focus order, the relevance of alt text, captions and cognitive load remain a **manual** audit (WCAG 2.1 AA). The passing landmarks (header/nav/main/footer), the skip-link and 66 aria-labels detected are all positive signs.

CT-03 : Fix the flagged issues, then commission a manual WCAG audit

Observation. Two binary audits fail: color-contrast (3 elements below the required ratio) and link-name (3 links with no discernible name — typically icon-only links). **Legal frame (stated precisely).** The **European Accessibility Act** has applied to consumer e-commerce in the EU since **28 June 2025**; your hreflang shows EU-facing storefronts (de-de, fr-fr, es-es, nl, be...), so this framework is directly relevant to your business. In the US, WCAG 2.1 AA is the de-facto reference (ADA Title III for private business). We frame this as compliance + inclusion + image SEO — not as asserting that a specific public deadline binds you. **Solution.** Fix the 6 machine-flagged issues (raise contrast to $\geq 4.5:1$ for text; add `aria-label` to icon links), then commission a manual WCAG 2.1 AA audit covering keyboard, focus order and alt relevance.

Profile: Frontend | **Seniority:** Confirmed | **Effort:** 2 man-days (the flagged fixes) + a separate manual audit engagement

Expected gain: a stronger compliance posture for your EU markets, a more inclusive store, and cleaner semantics that also help crawlers and AI.

CT-11 : Complete image alt coverage & fix decorative SVG labels

Observation. 19 of 42 homepage `<img>` carry an alt attribute (Lighthouse's `image-alt` did *not* fail, so many of the rest are likely decorative, lazy-loaded or `srcset` sources — to classify, not to assume). Two inline SVGs use non-descriptive `<title>` values ("Untitled-2", "wamlogo"). **Solution.** Classify each image: decorative $\rightarrow$ `alt=""` (an empty alt is correct, not a violation); informative $\rightarrow$ a concise descriptive alt (which doubles as image-SEO text). Give decorative logo SVGs `aria-hidden="true"` or a meaningful `<title>`.

Profile: Frontend | **Seniority:** Junior | **Effort:** 1.5 man-days

Expected gain: better screen-reader experience *and* more indexable image context (Google Images, AI).

International — a genuine strength to protect

Many brands that call themselves "worldwide" ship a single English storefront with a Google Translate widget. You do not: you run a real multi-region setup on Shopify Markets.

25

HREFLANG LOCALES

4

LANGUAGES (EN/DE/FR/ES)

x-default

FALLBACK DECLARED

The declared cluster spans en-us, en-ca, en-gb and a full European matrix (de-de, de-at, fr-fr, fr-be, es-es, nl, lu...), with a country selector (Orbe) and a proper x-default. This is indexable, per-language internationalisation — the kind Google can rank market by market.

Keep it clean. The main hygiene actions here are light: add og:locale:alternate (CT-07) and confirm the reciprocity of the hreflang cluster in Search Console. No re-platforming needed — the hard part is already done.

Because you already sell into the EU, this strength is also what makes the accessibility framework (CT-03) directly relevant — the two sections reinforce each other.

Competitive benchmark — where the market bar sits

We measured two named competitors with the **same** tool as your store — the PageSpeed Insights accessibility score (Google's own Lighthouse) — and crawled their public SEO capabilities directly. Other market references named in 2026 coverage: Sony, Anker / Soundcore, Bose, Marshall, Urbanista.

Layer 1 — your measured facts vs. the 2026 standard

The bar the market now rewards is valid, rich structured data and an accessible, international storefront. On the tested pages you meet most of it — but your AggregateRating is null, so you are **ineligible today for the review-star rich result** that a correctly-marked-up competitor can earn on the same query. Your accessibility (93) is at the bar, not ahead of it — and the EU market you serve raises that bar.

Layer 3 — measured comparison (homogeneous)

Measured / observed axis	Example Shop	JBL	Beats	Source
Accessibility (PSI mobile a11y)	93/100	94/100	92/100	PSI accessibility
Home JSON-LD types	Organization, WebSite	none detected	Organization, ContactPoint	crawl
hreflang links (home)	25	0	28	crawl
Sitemap.xml reachable	200	403	200	HTTP
Agentic llms.txt	200	403	404	HTTP

Homogeneous measure: PSI accessibility is Google's own Lighthouse, run identically on all three. SEO capabilities (structured data, hreflang, sitemap, llms.txt) are observed on both sides. Accessibility figures are single-run PSI lab values

— deterministic for accessibility audits, but not a performance benchmark (performance is out of scope for this discoverability audit).

Layer 2 — capability gaps (what an audio buyer / agent sees)

Capability the market expects	You	Named competitor	Evidence
Valid review stars in SERP (AggregateRating)	Null rating — ineligible	Achievable with correct markup	PDP \$structured data
Breadcrumb rich result	No BreadcrumbList	Common on mature catalogues	PDP / collection
Contactable brand entity (ContactPoint)	Absent	Beats: present	home JSON-LD
Indexable international (hreflang)	25 locales	Beats: 28 · JBL: 0 (home)	crawl
Agentic-commerce discovery layer	Live (Shopify)	JBL / Beats: not exposed	llms.txt / UCP

Layer 4 — where you are ahead, told straight. You are at parity on accessibility, ahead of JBL on home structured data, hreflang and an accessible sitemap, and you lead both competitors on the agentic-discovery layer (because Shopify shipped it). That is good news — but speed of plumbing is not the finish line. Buyers and AI agents decide on the *details*: valid review stars, breadcrumbs, a contactable entity. Beats already exposes a `ContactPoint` you do not; and until your `AggregateRating` is valid, a competitor with correct markup wins the star-rated listing on shared queries. The gap is narrow and specific — which is exactly why it is worth closing now, before the traffic peak.

Bottom line: on the criteria the market now treats as table stakes, your shortfall is concentrated in a few high-leverage details (valid ratings, breadcrumbs, entity contact data) rather than spread across the whole stack — and every one of them is a fast, native fix.

Timing — why now

We are 18 July 2026. For a US audio brand with a strong European footprint, the discoverability-relevant windows ahead are:




Amazon Prime Day sits around now (mid-July); the next owned-traffic surge is **Back-to-School** (~3-7 weeks out) — prime season for earbuds, gaming headsets and student gifting — followed by the Q4 peak of **Black Friday** → **Cyber Monday** → **Christmas**, the biggest audio-gifting window of the year.

The crawl-lag argument. Structured-data changes (valid ratings, breadcrumbs) take *weeks* to be re-crawled and reflected as rich results, and AI engines refresh their view on their own cadence. Fixing CT-01 and CT-02 **now** positions you to own star-rated listings and AI answers *before* Back-to-School and Q4 traffic arrives — not during it, when it is too late to influence how you appear.

Projection, honestly. We do not have your analytics, so we do not invent a percentage. The gain from review stars is best expressed as a formula on your own numbers: *monthly gain* $\approx$ *product-page impressions* $\times$ Δ *click-through (from stars)* $\times$ *conversion rate* $\times$ *average order value*. A 30-minute scoping workshop with your Search Console + Analytics turns this into a concrete figure for your catalogue.

Effort distribution by profile


Total backlog identified: **≈ 13 man-days** of build work (plus a separate external manual WCAG audit engagement). The chart shows how that build effort splits across profiles.




A discoverability program dominated by SEO/content and front-end semantics — no heavy back-end build, which is what makes it fast.

Action plan — backlog & priorities


Phase 1 — quick wins (this month, ahead of Back-to-School)



Phase 2 — foundations (before Q4 peak)



Phase 3 — prospective (GEO / agentic, after the measurable work)



ID	Action	Profile	Effort	Phase
CT-01	Fix null AggregateRating on PDP (valid review stars)	Dev / SEO	1.5 md	1
CT-02	Add BreadcrumbList structured data	Dev / SEO	1 md	1
CT-04	Trim over-long PDP title	SEO	1 md	1
CT-06	Enrich Organization (ContactPoint, clean sameAs)	SEO	0.5 md	1
CT-07	Complete Twitter/X card + og:locale:alternate	Frontend	0.5 md	1
CT-12	Dedupe robots Sitemap line; submit child sitemaps	SEO / DevOps	0.5 md	1
CT-03	Fix color-contrast x3 + link-name x3; manual WCAG audit	Frontend	2 md + audit	2
CT-05	One semantic h1 per page	Frontend	1 md	2
CT-11	Image alt coverage audit + SVG labels	Frontend	1.5 md	2
CT-08	Collection ItemList / CollectionPage schema	Dev / SEO	1 md	2
CT-09	Customise agentic files via Liquid (prospective)	SEO / Content	2 md	3
CT-10	Document AI-crawler policy (prospective)	SEO	0.5 md	3

≈ 13 man-days

FULL BACKLOG (ALL 12 ACTIONS)

≈ 9 man-days

PRIORITISED ROADMAP (PHASE 1–2, EXCL. EXTERNAL AUDIT & PROSPECTIVE)

Nothing is dropped: the prospective GEO items (CT-09, CT-10) stay in the backlog with their effort — they simply sit behind the measurable, higher-certainty work.

In six months you can be here: product pages that earn review stars and breadcrumbs in Google, a contactable brand entity recognised across AI answer engines, an accessibility posture aligned with your EU markets, and agentic files that tell *your* story — all built on native Shopify levers, with no paid app and no re-platforming.

Preliminary external evaluation — publicly observable signals only, no internal access. AI-assisted workflow with manual validation of retained findings.

Author: Yassin BELHAJ SALAH, E-commerce Solutions Architect · © Yassin BELHAJ SALAH — 2026. All rights reserved.