

Example Shop

External Web & E-commerce Diagnostic
Performance · Security · Discoverability · Conversion

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

Storefront tested: US (/en_us/) | **Platform:** Headless Adobe Commerce (Magento
2) + Next.js

Date: 18 July 2026

Author: Yassin BELHAJ SALAH

Architect & Business Practitioner Magento / Adobe Commerce

Confidentiality: internal use — Example Shop


© Yassin BELHAJ SALAH — 2026

Executive Summary

Scope & method. This is an **external, preliminary evaluation** based only on publicly observable signals — no access to your code, admin, infrastructure, analytics or monitoring. It is produced by an **AI-assisted workflow with manual validation** of the findings retained here. Every figure below is either measured, attributed to a named source, or expressed as a formula to be completed with your own data — nothing is invented.

Measurement note — bot protection. example-shop.com is protected by **DataDome** (a strength for security). It returns HTTP 403 to automated clients, so direct crawling of the HTML, robots.txt and sitemap was not possible. To measure you fairly and honestly, all page metrics were obtained **Google-side via PageSpeed Insights** (lab + real-user CrUX field data) and via infrastructure probes (TLS, DNS, HTTP headers). Where the block prevents observation, this report says "*to confirm*" — never "*absent*".

Overall posture (measured)



What is already strong

- **Modern headless stack** — Adobe Commerce (Magento 2) with a Next.js storefront, delivered on Fastly CDN.
- **SEO fundamentals 100/100** on the tested home page (title, meta, canonical, hreflang, crawlable).
- **Real-user LCP is FAST** and **CLS is FAST** (CrUX field) — your visitors see content quickly and the layout is stable.


Where the upside is

- **Interactivity (INP)** is your real-user gap — field rated AVERAGE while the market leader is FAST across all Core Web Vitals.
- **JavaScript weight** — 3.26 s bootup, 350 KiB unused JS, plus a 754 ms entry redirect (lab).
- **Accessibility of the underlying code** — real Lighthouse failures remain behind the overlay, at a moment when the EU Accessibility Act is in force.

- **Images already optimised** — Lighthouse finds 0 KiB recoverable on modern formats / offscreen images.
- **Strong email & transport security** — DMARC p=reject, HSTS with preload, TLS 1.3, DataDome, OneTrust consent.
- **Premium discovery engine** — Constructor.io AI search & New Relic real-user monitoring detected.
- **Response security headers** — Content-Security-Policy and hardening headers not observed on tested responses (Observatory D-).
- **Structured data & AI-search (GEO)** — to validate and enrich as AI answer engines reshape discovery.

Maturity radar (indicative, /10)

Synthesis score per axis from the measured signals in this report. It is a directional read to frame priorities, not a certification.



Performance 5 · SEO 7 · Security 5 · UX/Accessibility 6 (/10)

Headline. Example Shop runs a genuinely modern commerce stack and your real users already get a fast, stable first paint. The value on the table is **consolidation**: shave the JavaScript and redirect cost to close the INP gap with the field leader, harden response-level security, and make the underlying accessibility and structured data as strong as your visuals — all before the autumn-to-winter revenue ramp.

Platform & Technical Profile

A composable, headless architecture

The signals captured from Google's render point to a modern, decoupled stack — a real asset that many outdoor brands are still migrating toward:

Layer	Detected signal	Reading
Commerce backend	/graphql, /media/attribute/swatch/...	Adobe Commerce (Magento 2), GraphQL-driven
Storefront	/_next/static/chunks/*.js	Next.js (React) headless frontend
CDN / edge	x-served-by: cache-...-PAR, x-cache	Fastly
Search & discovery	ac.competitor-b.example	Constructor.io — AI-powered search (detected)
Bot protection	x-datadome: protected	DataDome (detected, in place)
Consent	cdn.competitor-c.example, competitor-d.example	OneTrust (detected)
Monitoring	bam.competitor-e.example, competitor-f.example	New Relic RUM (detected)
Tag management	analytics.example-shop.com/gtm.js	Google Tag Manager, first-party proxied (detected)
Integrator signal	security-hub.vaimo.network	Vaimo (Adobe/Magento partner)

Every third party above is a *detected asset/host* from the rendered page — a signal, not proof of active runtime configuration. Historical context (attributed): Vaimo delivered Example Shop's Adobe Commerce programme; public sources cite ~55 stores, 10 languages and 65,000+ products per store (*Vaimo case study; Essential Retail*) — figures to confirm against your current setup, not measured here.

Why this matters commercially. A headless Adobe Commerce + Next.js stack gives you a fast, flexible storefront and a deep commerce engine underneath. The trade-off of any React storefront is **JavaScript cost** — which is exactly where the measured performance headroom sits (next section). The good news: the fix is engineering discipline on an architecture you already own, not a re-platform.

Certificate & domains observed

Item	Observed	Status
TLS versions	TLS 1.2 + TLS 1.3	Modern
Certificate	Let's Encrypt, valid 22 Jun → 20 Sep 2026	Valid
Covered hosts (SAN)	example-shop.com, www, shop, commerce, competitor-g.example, www.hhworkwear	Consolidated

The certificate reveals a distinct **professional / workwear** presence (competitor-g.example) alongside the consumer store — a strategic thread picked up in the roadmap.

Native Platform Capabilities — Unexploited Levers

Adobe Commerce ships a large set of merchandising and loyalty capabilities **at no extra licence cost** beyond your edition. Because the storefront is bot-protected, we could not confirm which are live from the outside — so each row below is a **lever to confirm against your configuration and edition**, with the business reason to switch it on. This is the fastest kind of value: activation, not development.


Native capability (Adobe Commerce)	Status	Business lever
Related / Cross-sell / Up-sell	● to confirm (bot-protected)	Raise average basket by surfacing complementary layers & accessories — measure with your data.
Catalog Price Rules	● to confirm	Automate seasonal promotions (end-of-season, Black Friday) without developer time.
Cart Price Rules	● to confirm	First-order incentive to capture emails; threshold free-shipping to lift basket.
Gift Cards	● to confirm	New payment tender and a natural gifting driver for the December peak.
Reward Points / Loyalty*	● to confirm per edition	Recurring engagement for a repeat-purchase, seasonal-gear audience.
Store Credit*	● to confirm per edition	Frictionless returns/exchanges — reduces refund leakage, keeps revenue in-house.
Multiple Wishlists	● to confirm	"Ski kit", "sailing kit" lists — intent capture ahead of the season.
Gift Messages / Gift Wrap	● to confirm	Premium gifting experience for Q4 — differentiates the checkout.
Store Pickup / Click & Collect	● to confirm	If retail/dealer network exists — bridges online and physical demand.

*Reward Points and Store Credit are native to specific Adobe Commerce editions — **to confirm per your licence/entitlement**, never assumed included. Only native or free capabilities are recommended; no paid extension or marketing platform is proposed anywhere in this report.

Constructor.io is a real edge. You already run a premium AI search & discovery engine — a capability several competitors do not expose. The opportunity (CT-11) is to exploit it more fully: AI-ranked category pages, recommendations and "complete the look", so the investment pays back across the whole journey, not just the search box.

Performance & Core Web Vitals

All values measured via PageSpeed Insights (Google), mobile, 18 July 2026. Lab = synthetic throttled test (cold load); Field = CrUX real-user data over 28 days.



Read the two numbers honestly. The lab LCP of 14.3 s looks alarming, but it is a cold, throttled first-view that also absorbs a 754 ms entry redirect. Your **real users tell a better story**: CrUX rates field LCP and CLS as **FAST**. The genuine real-user weakness is **INP (interactivity), rated AVERAGE** — driven by JavaScript execution, not by images. So the work is targeted, not a rescue.

What the measurement shows

Metric	Measured	Reading
Total transfer	3,291 KiB (~3.3 MB)	Heavy — mostly JS
Requests	247	High
JS bootup (execution)	3,262 ms	Main INP contributor
Main-thread work	4,334 ms	High
Unused JavaScript	350 KiB / ~1,050 ms	Recoverable
Entry redirect /en_us/ → /en_us	754 ms (lab)	Avoidable
TBT (lab)	670 ms	needs-improvement
Modern images / offscreen	0 KiB recoverable	Already optimised ✓
Render-blocking resources	0 ms	Clean ✓

CT-01 : Cut JavaScript execution cost to close the INP gap

Observation: 3.26 s of JS bootup, 4.33 s of main-thread work and 350 KiB of unused JavaScript (≈1,050 ms) on the tested home page — the classic profile of a React storefront, and the direct driver of your AVERAGE field INP. **Risk:** interactivity lag on taps/filters costs conversions exactly when traffic spikes (Black Friday, winter peak). **Solution:** audit the Next.js bundle for route-level code-splitting, tree-shaking and dynamic imports; defer non-critical third parties (below); validate the appropriate strategy against your real build pipeline rather than applying a generic recipe.

Profile: Frontend | **Seniority:** Senior | **Effort:** 8 man/days

Expected gain: Lighthouse quantifies ~350 KiB / ~1,050 ms of directly recoverable JS on the tested page; the target is to move field INP from AVERAGE toward the FAST band the market leader holds (measured after deployment via your New Relic/CrUX field data).

CT-02 : Remove the entry redirect

Observation: the tested entry URL /en_us/ redirects to /en_us, costing 754 ms in the lab run. **Risk:** a redirect on the first navigation delays first paint and wastes crawl budget. **Solution:** normalise trailing-slash handling at the Fastly edge / routing layer so the canonical URL is served directly; validate canonical entry points across locales.

Profile: DevOps | **Seniority:** Confirmed | **Effort:** 1 man/day

Expected gain: -754 ms measured on the tested entry path (lab).

CT-03 : Govern third-party JavaScript

Observation: the render pulls GTM, Google Ads, Constructor.io, New Relic, OneTrust, EqualWeb and obs.competitor-h.example among others. Each adds main-thread work. **Risk:** third-party scripts are the least-controlled part of INP/TBT and can regress silently. **Solution:** build a tag inventory with per-vendor main-thread cost (from Lighthouse traces), consent-gate and lazy-load non-critical tags, and set a performance budget in CI. Treat this as a standing discipline, not a one-off.

Profile: Frontend | **Seniority:** Senior | **Effort:** 4 man/days

Expected gain: reduced main-thread contention (measured per-vendor); fewer INP regressions during campaign spikes.

CT-12 : Validate edge cache & TTFB

Observation: tested responses returned x-cache: MISS on Fastly. This is a single external sample on a bot-challenged path and does *not* prove a cache problem. **Risk:** under-cached HTML inflates TTFB at peak. **Solution:** with infrastructure access, review the Fastly / Varnish / Full-Page-Cache hit ratio and cacheability rules for storefront routes — a hypothesis to validate internally, not a confirmed defect.

Profile: DevOps | **Seniority:** Senior | **Effort:** 3 man/days

Expected gain: to be quantified once cache-hit ratio is measured with internal access.

Competitive Benchmark

Competitors were measured with the **same method as Example Shop** — PageSpeed Insights (Google), mobile, lab + CrUX field — so the comparison is homogeneous and legitimate. Named references in your

space: **The North Face, Patagonia, Arc'teryx** (broad technical outdoor); adjacent specialists include Norrøna, Columbia, Musto and Gill.

Layer 1 — your measured numbers vs the live market standard

Axis (measured)	Example Shop	The North Face	Patagonia	Arc'teryx	Source
Perf score (lab)	48	29	49	33	PSI lab (Google)
LCP (lab)	14.3s (poor)	17.8s	8.0s	7.1s	PSI lab
TBT (lab)	670 ms	1,570 ms	320 ms	1,540 ms	PSI lab
CrUX overall (real users)	AVERAGE	SLOW	FAST	AVERAGE	PSI field
CrUX LCP (real users)	FAST	FAST	FAST	FAST	PSI field
CrUX INP (real users)	AVERAGE	SLOW	FAST	AVERAGE	PSI field

Honest read. On synthetic lab score, Example Shop (48) ranks ahead of The North Face (29) and Arc'teryx (33), and is level with Patagonia (49). On real-user field data, your LCP is already FAST — a genuine, measured strength. **But Patagonia is the bar:** it is the only reference here rated **FAST across every Core Web Vital**, while you sit at AVERAGE, held back by INP. The market standard the leader has set is "fast everywhere" — and the distance between you and them is precisely the JavaScript/interactivity work in CT-01/CT-03.

Layer 2 — observable capability gaps

Capability the market expects	You	Reference	Evidence
AI search & discovery	Yes — Constructor.io	Uneven across peers	signature competitor-b.example
Fast-everywhere real-user CWV	INP AVERAGE	Patagonia FAST	CrUX field
Response security headers (CSP)	Not observed	Market baseline rising	Observatory D-
Code-level accessibility (EAA)	Failures behind overlay	EU baseline since 06/2025	Lighthouse a11y 89
AI-search citability (GEO)	to verify	Emerging standard	structured data / llms.txt

Section conclusion. Where the market now treats speed-everywhere and hardened, accessible storefronts as table stakes, Example Shop leads some peers but trails the field leader on interactivity and lags an emerging baseline on response security and code-level accessibility. Each gap is a reason an equally-equipped buyer could choose a competitor — and each is measured, not assumed.

Measurement note: all four competitors returned PSI data; no competitor score is estimated or invented. Comparison rests on your real measurements, the live Core Web Vitals standard, and capabilities observed on both sides.

Security

Assessed from the outside only. Where a conclusion would need internal access, it is stated as a hypothesis to validate — never as a confirmed internal defect.

Strong foundations (acknowledged)

Control	Observed	Status
Transport (TLS)	TLS 1.2 + 1.3	Modern ✓
HSTS	max-age=31536000; includeSubDomains; preload	Strong ✓ (already preloaded)
Email — DMARC	p=reject (Proofpoint reporting)	Strong ✓
Email — SPF	present (~all)	Present ✓
Bot protection	DataDome	In place ✓
Consent management	OneTrust	In place ✓

You are ahead here. HSTS is already preloaded and DMARC is at p=reject — the strongest settings, which many retailers never reach. The level above is operational: **monitor the DMARC aggregate reports** you already collect, and evaluate **MTA-STS** and **BIMI** as the next maturity step for brand-protected email.

CT-04 : Introduce a Content-Security-Policy

Observation: Mozilla Observatory grades the host **D- (25/100)** and Lighthouse best-practices is **73/100**. On the responses we could test (the bot-challenge response), only HSTS was present — no Content-Security-Policy. **Risk:** without a CSP, a single injected/third-party script has fewer guardrails — material given the number of tags in play. **Solution:** define a policy in `report-only` first, tune against real third-party origins, then enforce. Validate on authenticated storefront responses, not just the challenge page.

```
# Example, report-only first – tune to your real third-party origins, then
enforce:
Content-Security-Policy-Report-Only: default-src 'self';
  script-src 'self' 'unsafe-inline' *.example-shop.com *.competitor-b.example
*.competitor-f.example
  *.datadome.co *.competitor-d.example www.competitor-i.example;
report-uri /csp-report;
```

Profile: DevOps | **Seniority:** Architect | **Effort:** 6 man/days

Expected gain: a measurable reduction of the script-injection attack surface; improved Observatory / best-practices signals (final grade measured after deployment, not projected here).

CT-05 : Complete the hardening headers

Observation: on tested responses, X-Frame-Options/frame-ancestors, X-Content-Type-Options, Referrer-Policy and Permissions-Policy were not observed. **Risk:** clickjacking, MIME-sniffing and referrer leakage exposure. **Solution:** add the baseline headers at the edge; confirm they are present on storefront and checkout responses (not just the home path we could reach).

```
add_header X-Content-Type-Options "nosniff" always;
add_header X-Frame-Options "SAMEORIGIN" always;
add_header Referrer-Policy "strict-origin-when-cross-origin" always;
add_header Permissions-Policy "geolocation=(self), camera=()" always;
# HSTS is already strong and preloaded – keep as-is; no change to
includeSubDomains/preload.
```

Profile: DevOps | **Seniority:** Confirmed | **Effort:** 2 man/days

Expected gain: closes common header gaps; hardening verifiable via Observatory re-scan after rollout.

Discoverability — SEO & AI Search (GEO)

Your SEO base is excellent. Lighthouse rates the tested home page **SEO 100/100**: valid title and meta description, canonical, hreflang present (critical for your 10-language footprint) and the page is crawlable. This is a strong platform to build on.

CT-06 : Validate & enrich structured data on product pages

Observation: Lighthouse treats structured data as a *manual* audit, and the bot protection prevented us from crawling a product page — so we cannot confirm the state of Product/Offer/AggregateRating/BreadcrumbList markup from the outside. **Risk:** incomplete product schema weakens rich results (price, availability, ratings) and, increasingly, AI-answer eligibility. **Solution:** validate each template on a live PDP with Google's Rich Results Test; complete any gaps in the module/template appropriate to your version. State findings per page tested — not site-wide — until every template is sampled.

Profile: SEO | **Seniority:** Confirmed | **Effort:** 4 man/days

Expected gain: eligibility for richer result types; a prerequisite for the GEO work below.


CT-07 : AI-search citability (GEO) — prospective

Context (dated veille). A growing share of discovery now happens inside answer engines (ChatGPT, Perplexity, Google AI Overviews). Zero-click searches on Google rose from 56% to 69% in the year after AI Overviews launched (*Search Engine Land*, "Mastering GEO in 2026"). Princeton's GEO research finds that citation-worthy elements — citing sources, statistics, quotations — can lift AI visibility by 30–40% (*Princeton GEO study*). **Observation:** we could not verify your robots.txt AI-crawler directives or the presence of an llms.txt because the file was bot-blocked — to check internally. **Solution:** review AI-crawler access policy deliberately, ensure key content is server-rendered (a headless-React risk), consider an llms.txt, and structure PDPs/guides with clear answers and rich schema.

Profile: SEO | **Seniority:** Confirmed | **Effort:** 3 man/days (prospective)

Expected gain: improved eligibility for AI-answer citation — a strategic bet, ranked below the measured performance/security/accessibility work (prioritisation is judgement, not a measurement). Sources: *Search Engine Land 2026*; *Princeton GEO research*; *LLMrefs 2026 GEO guide*; *Google AI-optimization guide*.

Accessibility & Compliance



CT-08 : Fix accessibility in the code — beyond the overlay

Observation. Lighthouse rates the tested home page **89/100** with named, code-level failures that an overlay does not resolve:

Lighthouse audit	Meaning
heading-order	Headings not in a logical sequence — hurts screen-reader navigation.
link-name	Links without a discernible accessible name (also flagged in SEO).
list	List markup not correctly structured.
target-size	Tap targets below the recommended size on mobile.

Why the overlay is not the answer. An **EqualWeb** accessibility overlay is detected on the site. Overlays are a genuine signal that accessibility is on your radar — but they inject JavaScript on top of the markup and **do not fix the underlying code** (the four failures above remain). The documented evidence is consistent: automated tooling and overlays catch only ~20–30% of WCAG issues, with 70–80% requiring manual testing (*Overlay Fact Sheet*); the U.S. FTC finalised a \$1M order against a leading overlay vendor in April 2025 over deceptive WCAG-compliance claims (*FTC, primary source*); and 25% of digital-accessibility lawsuits in 2024 targeted sites that *used* overlays (*UsableNet 2024*). The European Commission's position is that overlays do not constitute a valid path to compliance.

Why now. The **European Accessibility Act is in force since 28 June 2025**, requiring EN 301 549 / WCAG 2.1 AA for e-commerce across the EU — directly relevant to your multi-market European footprint. This report names no specific case against Example Shop; it states the general, documented framework and the measured gaps on your own pages.

Solution. Remediate at source in the Next.js components (heading hierarchy, accessible link/button names, semantic lists, mobile target sizing), then commission a manual WCAG 2.1 AA audit with assistive-technology testing. Keep or retire the overlay as a convenience layer — but do not rely on it for compliance.

Profile: Frontend | **Seniority:** Senior | **Effort:** 10 man/days (code) + manual audit

Expected gain: genuine WCAG 2.1 AA progress and reduced legal exposure under the EAA; re-measured via Lighthouse + manual audit after remediation (no post-fix score projected here).

CT-09 : Front-end best-practices hygiene

Observation: Lighthouse best-practices is **73/100**, with deprecations (deprecated browser APIs in use) and valid-source-maps (source maps shipped/exposed in production). **Risk:** deprecated APIs break over time; exposed source maps reveal application internals. **Solution:** remove deprecated calls flagged by Lighthouse and stop serving production source maps (or restrict them) in the Next.js build.

Profile: Frontend | **Seniority:** Confirmed | **Effort:** 3 man/days

Expected gain: cleaner best-practices signal; smaller information-leak surface.

Conversion & Merchandising

The storefront is bot-protected, so runtime conversion features could not be exercised directly. The recommendations below build on capabilities you already own (Adobe Commerce native, Constructor.io) — no paid platform is proposed.

CT-10 : Activate native Adobe Commerce merchandising

Observation: the native merchandising levers in the capability table (cross-sell, price rules, gift cards, wishlists, gift messaging, click & collect) could not be confirmed live from the outside. **Risk:** leaving zero-cost basket and gifting levers idle through the Q4 peak. **Solution:** in a short workshop, confirm which are active in **Admin** → **Catalog / Marketing / Stores** and switch on the seasonal ones; quantify the basket lift with your analytics.

Profile: Dev | **Seniority:** Confirmed | **Effort:** 6 man/days

Expected gain: $\text{Gain} = \text{traffic} \times \Delta \text{conversion} \times \text{average basket}$ — to be quantified from your Analytics in a 30-minute scoping workshop.

CT-11 : Exploit Constructor.io & native recommendations end-to-end

Observation: a premium AI search engine (Constructor.io) is detected but its reach beyond the search box is unknown from the outside. **Risk:** paying for AI discovery while category, PDP and cart recommendations run on defaults. **Solution:** extend AI ranking to category pages, "complete the look" and cart cross-sell; A/B test recommendation slots. **Solution** pairs naturally with CT-10's native cross-sell.

Profile: Dev | **Seniority:** Senior | **Effort:** 7 man/days

Expected gain: higher discovery-to-cart rate — measured via on-site A/B testing with your data.

CT-13 : Unify the consumer & professional experience — strategic

Observation: the certificate reveals a separate **professional / workwear** presence (competitor - `g.example.commerce.example-shop.com`) distinct from the consumer store. **Opportunity:** a shared design system, account and content layer across both audiences on the same Adobe Commerce foundation can reduce duplication and open a cleaner professional buying journey. **Solution:** a discovery track to map the two experiences and a convergence roadmap — capabilities and entitlements to confirm against your Adobe Commerce licence.

Profile: Fullstack | **Seniority:** Senior | **Effort:** 8 man/days (prospective)

Expected gain: reduced maintenance duplication and a stronger professional journey — to be sized in discovery.

CT-14 : Wire Core Web Vitals field alerting into New Relic


Observation: New Relic RUM is already detected — you have real-user telemetry. **Opportunity:** put INP/LCP/CLS field thresholds on a dashboard with alerting, so the CT-01/CT-03 gains are tracked and protected. **Solution:** add Web-Vitals field metrics and alert on INP regressions, especially around campaign launches.

Profile: DevOps | **Seniority:** Confirmed | **Effort:** 2 man/days

Expected gain: continuous, measured guardrail on the interactivity work.

Business Timing & Seasonality

As of July 2026, we are in the Northern-hemisphere summer — peak sailing season and, more importantly for a technical-outerwear brand, the **start of the autumn-to-winter buying ramp**. The commercial calendar ahead is dense:



For Example Shop, the December–February ski & winter window is the highest-revenue season, immediately preceded by Black Friday. That places us **roughly four months from the most important traffic surge of the year**.

Why the timing works in your favour. The measured headroom — JavaScript/INP (CT-01/03), the entry redirect (CT-02), security headers (CT-04/05) and accessibility (CT-08) — is exactly the foundation you want *solid before* traffic peaks, not during. Starting now means the quick wins land before the ramp and the larger items before Black Friday.

How to value the gains (no invented numbers)


- **Measured:** ~350 KiB / ~1,050 ms recoverable JS (CT-01) and ~754 ms on the tested entry redirect (CT-02) — Lighthouse figures for today, not a projected end-state.
- **Attributed benchmark:** the live Core Web Vitals standard rates "fast-everywhere" as the leader's bar (Patagonia, FAST field) — a market reference, not a measurement of your revenue.

- **Client formula:** Monthly value = traffic × Δ conversion rate × average basket. Your Analytics let us quantify this in a 30-minute scoping workshop — the natural next step.

You already have strong fundamentals — SEO 100, fast real-user LCP, a modern stack and premium search. Prioritising the three quick wins first lets you feel an impact before the winter peak.

Effort Distribution by Profile

Full identified backlog: **67 man/days** across 14 workstreams. The chart shows how that effort splits by skill profile.




Prioritised Roadmap


Sequenced across three phases to land quick wins and the security/accessibility foundation **before the autumn-to-winter ramp**, then growth. The full backlog is listed — nothing is dropped; lower-priority items simply move to Phase 3.

#	Workstream	Profile	Effort	Phase
CT-02	Remove entry redirect (–754 ms)	DevOps	1 j	1
CT-05	Complete hardening headers	DevOps	2 j	1
CT-06	Validate/enrich PDP structured data	SEO	4 j	1

CT-09	Best-practices hygiene (deprecations/source maps)	Frontend	3 j	1
CT-14	CWV field alerting in New Relic	DevOps	2 j	1
CT-01	Cut JS execution / close INP gap	Frontend	8 j	2
CT-03	Third-party script governance	Frontend	4 j	2
CT-04	Content-Security-Policy (report → enforce)	DevOps	6 j	2
CT-08	Accessibility code remediation (EAA/WCAG)	Frontend	10 j	2
CT-10	Activate native merchandising levers	Dev	6 j	3
CT-11	Exploit Constructor.io end-to-end	Dev	7 j	3
CT-12	Validate edge cache & TTFB	DevOps	3 j	3
CT-07	AI-search citability (GEO) — prospective	SEO	3 j	3
CT-13	Unify consumer & professional experience — prospective	Fullstack	8 j	3

Timeline (≈6 months)





Six-month projection. By the winter peak you can have a leaner, more interactive storefront (closing the INP gap with the field leader), hardened response security, genuinely accessible code aligned with the EAA, and richer product data feeding both classic search and AI answers — built entirely on the modern stack you already run.

Example Shop — External Web & E-commerce Diagnostic · Prepared by Yassin BELHAJ SALAH, Architect & Business Practitioner Magento / Adobe Commerce · 18 July 2026

External preliminary evaluation from publicly observable signals; AI-assisted workflow with manual validation.
Measurements: PageSpeed Insights (Google, lab + CrUX field), Mozilla Observatory, TLS/DNS/HTTP probes — 18 July 2026.

© Yassin BELHAJ SALAH — 2026. All rights reserved.