

TECHNICAL AUDIT · PERFORMANCE · SEO

SMART-UI.PRO

AUDIT REPORT

A full technical read of one production website: Core Web Vitals, rendering, third-party cost, crawlability, structured data and security headers. Every number in this document was measured, not estimated.

SAMPLE REPORT

THE VERDICT

Perfect on desktop. Slow on mobile.

The infrastructure is right: the server answers in **30 ms**, there is **zero layout shift**, and the main thread is essentially never blocked. On desktop the site scores a clean **100/100** and paints its largest element in **0.8 s**.

On a mid-range phone the same page takes **5.6 s** to paint that element — more than double Google's 2.5 s threshold. Nothing is broken; the page simply spends its mobile budget on **409 KB of JavaScript** and on three third-party origins it connects to too late.



Same page, two devices



METRIC	MOBILE	DESKTOP	GOOGLE THRESHOLD
Largest Contentful Paint	5.6 s	0.8 s	≤ 2.5 s
First Contentful Paint	3.1 s	0.4 s	≤ 1.8 s
Speed Index	4.4 s	0.6 s	≤ 3.4 s
Total Blocking Time	30 ms	0 ms	≤ 200 ms
Cumulative Layout Shift	0	0	≤ 0.1
Server response	30 ms	30 ms	≤ 600 ms

Why this matters commercially. The people who build and review a site look at it on a desktop, where it is flawless. The people who buy from it increasingly do not. Mobile LCP is the metric Google uses for ranking, and 5.6 s is the difference between a page that feels instant and one a visitor abandons before it renders.

HOW THIS WAS MEASURED

Every claim is reproducible

Audits that are written without touching the site read plausibly and age badly. This one was produced entirely from live measurements, and each finding carries the evidence that produced it.



Reproduce the headline number

```
$ npx lighthouse@12 https://www.smart-ui.pro/ \
  --form-factor=mobile --screenEmulation.mobile \
  --only-categories=performance --output=json

✓ largest-contentful-paint    5.6 s
✓ total-blocking-time        30 ms
✓ cumulative-layout-shift    0
performance score 71/100 · measured 24 Aug 2026
```

Two findings we deleted

A first pass flagged **invalid structured data** and a **canonical mismatch** on the English URLs. Both were checked before they were written down, and both were wrong:

- The "invalid JSON-LD" was the Next.js RSC payload (`self.__next_f.push`) caught by a careless regex. All three real blocks — Organization, BreadcrumbList, Person — parse cleanly.
- The "canonical mismatch" on `/en/about-us` was not one: that URL answers 301, and the canonical correctly names the destination.

We report what survives verification. A finding you cannot reproduce is a finding that costs your team a day and returns nothing.

PRIORITISED FINDINGS

What to fix, in order

FINDING 01

P1 · CRITICAL

Mobile LCP is 5.6 s against a 2.5 s threshold

The largest element is the hero image, served through `/_next/image` and already marked `fetchpriority="high"`. It paints in 0.8 s on desktop, so the image is not the problem — the mobile budget is spent before the browser reaches it.

EVIDENCE

Lighthouse mobile preset. The LCP element is the hero image in both profiles: desktop 0.8 s, mobile 5.6 s, identical markup.

FIX

Preload the hero source, cut the JavaScript ahead of it (findings 02 and 03), then re-measure. Expect 2.0–2.5 s without touching the design.

FINDING 02

P1 · QUICK WIN

Three third-party origins are connected too late

The browser discovers `google.es`, `stats.g.doubleclick.net` and `region1.analytics.google.com` only when it needs them, paying full DNS, TCP and TLS cost on the critical path.

COST

318 MS · 318 MS · 317 MS

— Lighthouse estimates 320 ms recoverable

FIX

Three preconnect link tags in the document head. Roughly ten minutes of work.

FINDING 03

P2 · HIGH

96 KiB of JavaScript is downloaded and never used

Google Tag Manager accounts for **74 KiB** of it on its own, with 22 KiB more in one first-party chunk.

```
googletagmanager.com/gtag/js?id=G-SGY5BK8163 74 KiB unused
_next/static/chunks/3024-0b26158f...js         22 KiB unused
```

FIX

Load analytics after first interaction or move it off the main thread; split the first-party chunk at its route boundary.

NOTE

This is a measurement, not an argument against analytics. It is an argument about *when* it loads.

FINDING 04

P2 · HIGH

Third parties are 30% of the page weight

9 of 58 requests come from other companies and carry **193 KB of 641 KB**. Scripts overall are 409 KB, 64% of the page, against 121 KB of images across 23 files.

EVIDENCE

Chrome resource timing, mobile profile

FIX

Set a third-party byte budget and enforce it in CI, the same way a test suite is enforced.

FINDING 05

P3 · SECURITY

No Content-Security-Policy header

Everything else on the transport layer is right: HSTS for two years, frame protection, MIME sniffing off, a sane referrer policy and a locked-down permissions policy. CSP is the one missing control, and it is the one that limits the damage of an injected script.

EVIDENCE

Response headers on
/

carry HSTS, X-Frame-Options, X-Content-Type-Options,
Referrer-Policy, Permissions-Policy — and no CSP

FIX

Ship

CONTENT-SECURITY-POLICY-REPORT-ONLY

first, collect violations for two weeks against the real third
parties, then enforce.

FINDING 06

P3 · SEO

A link that says only "Learn More"

Anchor text is a ranking signal and a screen-reader landmark. One link on the home page points at /ai-automation carrying no information at all.

FIX

Name the destination: "Explore AI automation". One line of copy.

FINDING 07

P3 · ACCESSIBILITY

Project titles fall below AA contrast

.project_name is painted #006EE0 on the dark section background and does not clear the 4.5:1 ratio WCAG AA requires for body text.

FIX

Lighten the blue or set it on a lighter surface. This is the only contrast failure on the page — accessibility already scores 95.

FINDING 08

P4 · POLISH

Three alt attributes repeat the visible caption

The project cards for San Lorenzo Foods, DuelMasters and Una Providore carry alt text identical to the caption printed beside them, so a screen reader announces each name twice. Separately, the hero image — the LCP element — carries an empty alt="", which marks the most prominent image on the page as decorative.

FIX

Empty the redundant attributes, and describe the hero properly.

WHAT WE ARE NOT GOING TO TOUCH

The part of the audit nobody sends you

An audit that only lists problems is easy to write and impossible to price. This is what was checked and found correct — so you do not pay anyone to fix it twice.

CHECKED	RESULT
Server response time	30 ms — twenty times inside the threshold
Cumulative Layout Shift	0 on both mobile and desktop. Nothing jumps.
Total Blocking Time	30 ms mobile, 0 ms desktop
Best Practices	100 / 100 — no console errors, no deprecated APIs, no insecure requests
Redirect chain	Single hop. http → https → www, each a clean 308
TLS	Let's Encrypt, valid to 15 Oct 2026 , HSTS max-age 2 years
robots.txt	Present and deliberate — explicit allow rules for GPTBot, ClaudeBot and PerplexityBot
sitemap.xml	24 URLs , well-formed, served at the canonical host
llms.txt	200 — ahead of most sites in this market
Structured data	Valid — Organization, BreadcrumbList, Person all parse
Image alt coverage	50 images on the home page, zero missing alt
Image weight	121 KB across 23 files. Images are not this site's problem — a claim worth testing before anyone proposes an image project
Canonicals	Correct , including on redirected language URLs

Read this section as a price signal. Twelve of the checks above needed no work. The engagement that follows this report is therefore small and specific — not a rebuild.

WHAT WE WOULD DO

Sequenced by return, not by severity

#	ACTION	EFFORT	EXPECTED EFFECT
01	Preconnect to the three analytics origins	10 min	–320 ms on the critical path
02	Defer Google Tag Manager past first interaction	2–4 h	–74 KiB of blocking JavaScript
03	Preload the hero image source	1 h	LCP starts earlier by roughly a second
04	Split the 3024 chunk at its route boundary	3–6 h	–22 KiB on first load
05	Re-measure and confirm mobile LCP under 2.5 s	1 h	Green Core Web Vitals
06	Fix anchor text, contrast and the four alt attributes	1 h	SEO 92 → 100, Accessibility 95 → 100
07	CSP in report-only, then enforce	1–2 wk	Last missing security control
08	Performance budget in CI	4 h	The result stops regressing silently

Total for items 01–06: roughly one working week. Items 07 and 08 are ongoing work, not a sprint. Nothing here requires a redesign, a migration or a new framework.

How we close it out

Every engagement ends in the same format it started: metric, before, after. Same tool, same device profile, same day of the week. If a number did not move, it says so.

```
$ smart-ai audit --url smart-ui.pro --compare baseline
LCP (mobile)      5.6s → target ≤ 2.5s
Unused JS         96 KiB → target ≤ 20 KiB
Third-party       193 KB → budgeted
SEO / A11y        92 / 95 → target 100 / 100
baseline captured 24 Aug 2026 · re-measured on delivery
```

HOW SLOW IS YOUR SITE, REALLY?

This is the report you get for free, in 48 hours, on your own domain — same tools, same format, same refusal to write down anything we have not measured.

NO OBLIGATION · ENGLISH OR SPANISH · REPORT IN 48H

EMAIL

hello@smart-ai.es

WEB

smart-ai.es

STUDIO

Mérida, España · CET