ARCHITECTURE NOTES · RESOURCE

Funding the Lighthouse-100 rebuild

A 92 on Lighthouse site converts measurably worse than a 100. The cost of closing the gap compounds across a year.

Companion to: The Lighthouse-100 rebuild every founder is afraid to fund.

Why this deck exists

Most founders treat 92 on Lighthouse as good enough. The article argues that on a real revenue page receiving meaningful traffic, the gap between 92 and 100 is measurable and the compounding cost across a year is non-trivial. This deck is for the founder weighing whether to fund a rebuild, the agency partner who needs to scope it, and the head of growth who keeps losing the budget conversation. It turns the architecture argument into a fundable scope, with the structural work named, the failure modes mapped, and the staging-versus-production gap made explicit.

- Why 92 is the local maximum agencies stop at
- The structural work that closes the last eight points
- How render-blocking JavaScript hides on the homepage demo
- Funding the rebuild as architecture, not styling

What the last eight points actually require

01 **Removing render-blocking JavaScript**

The marketing team installed analytics, attribution, and tag managers over time. Each one renders before paint. Removing them requires a conversation about which signals are load-bearing, which is rarely a styling decision.

02 **Deferring third-party tags**

Tags that need to fire eventually but not before paint. Defer, lazy-load, or replace with server-side equivalents. Each tag is a negotiation with the team that owns the signal it produces.

03 **Fixing layout shift in heroes**

Heroes that look fine on staging and break on real devices. The fix is reserving space, sizing media, and testing on the device profile that matches actual traffic, not the agency M2 Pro.

04 **Image format pipeline**

A real format pipeline. Modern formats served conditionally, sized correctly per breakpoint, with fallback rules that do not regress on older devices.

05 **Real-device testing**

The redesign demo runs in Chrome on the agency M2 Pro. Real users run on midrange phones with throttled connections. The test profile must match the user profile or the score is fiction.

Why 92 looks like 100 on the demo

Demo conditions (agency M2 Pro, Chrome)

- Fast CPU, high memory headroom
- Fiber connection, low latency
- Cached assets after first load
- No competing browser tabs
- Lighthouse run on the homepage only

Real-user conditions

- Midrange phone, thermal throttling
- Mobile network, variable latency
- Cold cache on first visit
- Heavy ad and tag stack loading
- Revenue pages, not the homepage

The compounding cost (illustrative, not a case study)

MONTHLY SESSIONS ON REVENUE PAGE

100,000

Meaningful traffic, the article qualifier

CONVERSION AT 92

baseline

The local maximum the redesign quotes to

CONVERSION AT 100

measurably higher

The article claim, framed as a real gap

ANNUAL COST OF THE GAP

trivial

Compounding across twelve months of traffic

Funding the rebuild as architecture work

- **Reframe the brief from styling to architecture**
  A redesign quoted to 92 funds a styling project. A rebuild quoted to 100 funds an architecture project. Different scopes.
- **Audit the full tag and script stack**
  Every render-blocking script gets a named owner and a deferral plan. No script is exempt from the conversation.
- **Set the test profile to match real users**
  Throttled CPU, throttled network, midrange device. The demo profile is not the user profile.
- **Score revenue pages, not the homepage**
  The homepage is the demo target. Revenue pages are where the conversion gap lives. Score the right page.
- **Reserve space for every async element**
  Layout shift in heroes is the most common cause of the last few points. Fixed dimensions, sized media, no surprises after paint.
- **Treat the score as a year-long metric**
  A 100 at launch that drifts to 92 in six months is the default trajectory. Budget the maintenance with the rebuild.
- **Name the owner of the score**
  A score with no owner is a score that drifts. Engineering leadership owns it, not marketing.

Anti-patterns we see in scopes

01 **Quoting to the agency demo profile**

The redesign quotes a number that gets the site to 92 because that is what the homepage demo looked like in Chrome on the agency M2 Pro. Real users on real devices see a slower site. The quote is honest about the demo, dishonest about the user.

02 **Treating the rebuild as styling**

When the rebuild is briefed as styling, the structural work falls out of scope. Render-blocking scripts stay, layout shift stays, the score stalls at 92. The fix is funding the rebuild as architecture work, not styling work.

03 **Calling 92 good enough**

On a hobby site, 92 is fine. On a real revenue page receiving meaningful traffic, 92 leaves measurable conversion on the table every day. The question is the traffic and revenue at stake, not the score in isolation.

04 **Skipping real-device testing**

A score from the agency laptop is a score for the agency laptop. Without real-device testing the team is optimizing for a user profile that does not visit the site.

Signals the rebuild is funded correctly

LIGHTHOUSE ON REVENUE PAGES

100

Not just the homepage

TEST PROFILE

real-device

Throttled CPU and network

RENDER-BLOCKING SCRIPTS

0

Or each one has a named owner and a defer plan

LAYOUT SHIFT

meaningful floo

Reserved space on every async element

NEXT STEP

Fund the rebuild as architecture work, not styling work.

A 92 on Lighthouse compounds against you across a year of meaningful traffic. The gap is structural and the structural work is what closes it.

Read the full architecture note ->