meticulosity

HubSpot Solutions Partner Program DIAMOND

TECHNICAL WHITE PAPER · AUGUST 2026

From HubSpot Content Hub to Next.js

An AI-executed site migration: how we rebuilt, rewrote and relaunched a 700-page site in days, and what happened to search, AI visibility and lead flow.

620	+5,335%	+1,010%	448
PAGES INDEXED, FROM ROUGHLY 40	IMPRESSIONS, 28 DAYS	CLICKS, 28 DAYS	BLOG POSTS REWRITTEN IN 19.5 HOURS

Published by
Meticulosity · The HubSpot Agency for Agencies

One month post-cutover
Measured 6 to 10 August 2026

WHAT IS INSIDE

Contents

Four parts: why the old site was dying, how the rebuild was executed, what the numbers did in the first month, and what we would tell another team. Every figure is a Google Search Console or GA4 capture, dated and sourced.

EXECUTIVE SUMMARY

A migration was the smallest part of what changed.

In July 2026 Meticulosity moved meticulosity.com off HubSpot Content Hub onto a statically generated Next.js site. The platform switch is not the story. The story is what a small team with AI-assisted engineering did **around** the migration, in hours rather than months.

620
PAGES INDEXED
From roughly 40 at the June trough

+5,335%
IMPRESSIONS, 28D
119K, www property

+1,010%
CLICKS, 28D
111 clicks in the window

2,060
GOOGLE AI IMPRESSIONS
From zero before 12 July

448
POSTS REWRITTEN
In roughly 19.5 hours

129
404 AFFECTED PAGES
Down from 951, still draining

$0
MEDIA INFERENCE COST
551 images, 459 narrations

99%
GTMETRIX HOMEPAGE
Up from 70% on the old stack

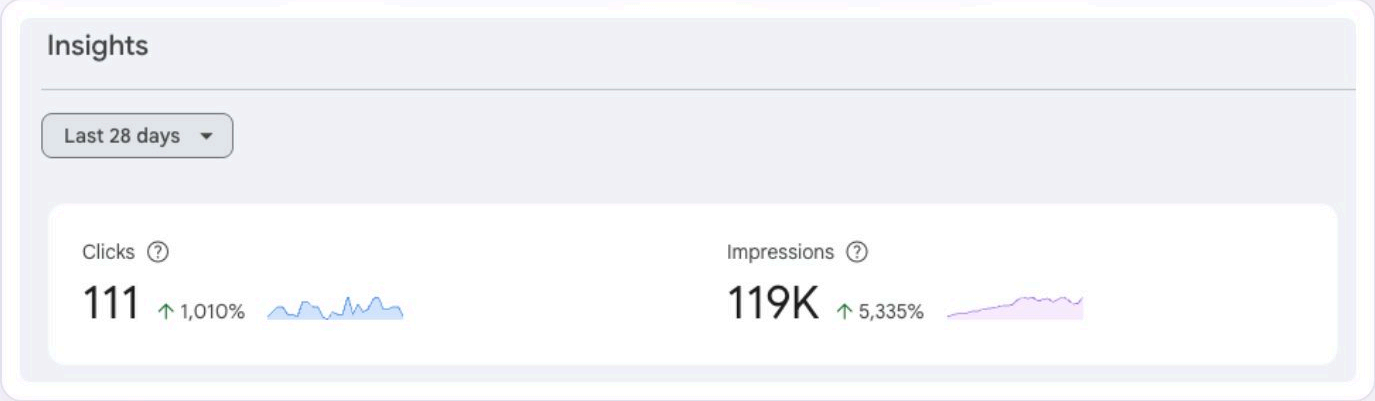

FIG 1 Search Console Insights, trailing 28 days at the one-month mark: 111 clicks and 119K impressions. Source: Google Search Console.

The counterintuitive part: this site's organic collapse **predated** the migration by months, and the migration is where every line turns back up. Part I tells that story honestly, because the lesson is bigger than a platform choice. A modern content stack only pays when you can also fix the content, at corpus scale, at AI speed.

PART ONE

Why we migrated

A Diamond HubSpot partner arguing itself off HubSpot Content Hub needs to show its working. Here is the measured baseline, and the honest post mortem of what actually killed the old site.

THE STARTING LINE

The stack we were on, and the measured baseline

meticulosity.com ran on HubSpot Content Hub: 54 marketing pages, a 551-post blog, HubSpot forms, HubSpot CTAs and HubSpot's platform URL space.

As a Diamond HubSpot partner we know the platform well, and this is not an anti-HubSpot argument. Content Hub is a fine CMS for teams that need marketing autonomy without engineering. Our problem was different: we needed **surgical, corpus-wide control** over rendering, structured data, URL space and performance, and we had AI-assisted engineering capacity that a closed CMS could not leverage.

Performance, four days before cutover

MEASURED	HOMEPAGE	MONEY PAGE
GTmetrix Performance	70%	69%
GTmetrix Structure	85%	91%
Largest Contentful Paint	1.4s	1.6s
Total Blocking Time	568ms	467ms
Time to First Byte	502ms	459ms

GTmetrix, 4 to 6 July 2026, Seattle test server, Chrome 142 / Lighthouse 12.6.1, Broadband Fast. Money page: /agency-hubspot-support; the homepage also carried 4.17MB across 208 requests. The archived final HubSpot build, re-measured on 10 August 2026, still scores mobile 70/100: the baseline is stable, not a bad day.

Search Console told a worse story

- **Indexed pages had drained to roughly 30 to 40**, from about 150 in May.
- **Google was crawling 50 to 100 requests per day** against roughly 5,400 known URLs: a full recrawl cycle measured in months.
- Clicks on the www property fell from 295 in March to **19 in April and 8 in June**. Even the brand query "meticulosity" went to zero clicks.
- The Coverage report carried **roughly 940 "Not found (404)" URLs**, 374 of them HubSpot platform bundle paths that Google kept recrawling on a schedule.
- Every blog post had an AMP twin and every CTA an encrypted redirect URL. At 50 to 100 crawls a day, Google was spending its residual attention on garbage.
- **The tell:** the old sitemap at one point listed only **31 URLs** of a roughly 670-URL site.

THE HONEST VERSION

What actually killed the old site

We reconstructed the decline from five years of data: Search Console, GA4, DataForSEO historical rankings and Semrush audit history. It runs in three phases, and the migration is not one of them.

1

The asset was mispriced from the start

The blog's search footprint was consumer how-to content, Facebook page management and clickbait explainers, ranking for queries no agency-services buyer ever types. Over the last 16 measurable months the 551-post blog produced **8,859 clicks against 15.37 million impressions**, a 0.06% click-through rate, and **29 tracked key events**. 42% of posts earned zero clicks.

2

In 2025 the SERP stopped paying for rankings we still held

Third-party rank tracking showed top-10 keyword counts at a multi-year high in November 2025 while clicks fell every single month. The Facebook cluster's impressions stayed flat all year while its clicks fell **74%**. AI Overviews and zero-click results absorbed exactly the query class the site ranked for.

3

Two Google core updates finished it

The December 2025 core update started the rankings slide. The March 2026 core update, 27 March to 8 April, lands exactly on the April cliff, with a portal migration churning the site in between (Semrush: errors 175 to 1,341, temporary redirects 266 to 2,005, broken internal links 112 to 629). Quality demotion collapses crawl scheduling, which drains the index, which at the extreme takes brand serving with it.

Two conclusions drove the project

First, there was no penalty to lift and nothing to wait out: **the only way back into the index was to become a different site.**

Second, becoming a different site meant the content, not just the platform. A lift-and-shift migration of the same corpus would have moved the same liability to faster hosting.

WHY THIS MATTERS TO YOU

If your own analytics show rankings holding while clicks fall, that is the same signal, and it is not a reporting artifact. It is the SERP re-pricing the query class you rank for. Read **clicks and conversions per query class**, never keyword counts.

THE STACK DECISION

Why Next.js specifically

Five properties, and only one of them is about speed. The decisive one is that the content model turns a corpus into something an agent can operate on.

Static generation

Every page pre-rendered at build time. No CMS runtime sits between Google and the HTML, and nothing can slow down under crawl load.

Content as code-adjacent data

One JSON file per page, one MDX file per post. The entire corpus becomes scriptable, diffable, verifiable in CI and, decisively, operable by AI agents at scale.

Total control of the head

Metadata, canonicals, Open Graph and every JSON-LD block generated per page from typed content, so the schema graph cannot drift away from what is on the page.

A URL space we own

Real 301s and real 410s from one config file, instead of platform bundles, AMP twins and encrypted CTA redirects we could neither see nor retire.

Performance as a design property, not a project

Self-hosted fonts, inline SVG, lazy third-party embeds and near-zero client JavaScript by default. The fast version is the default version, so nobody has to defend a performance budget in a sprint planning meeting six months from now.

THE ONE THAT MATTERED MOST

Of the five, **content as typed data** is the one that made everything in Part II possible. A closed CMS makes 448 rewrites, 551 image regenerations and 459 narrations into three separate projects. Typed files on disk make them three scripts.

What the new stack had to be able to survive.

755	687	557	105
BUILT PAGES	SITEMAP URLS	BLOG POSTS	MARKETING PAGES

II

PART TWO

How we did it

Everything that follows was executed with Claude Code driving the work: overnight unattended builds, parallel agent fan-outs for per-page and per-post work, and scripted verification gates between every stage. Dates and durations come from the git history.

PHASE 1 · 2 JULY

The overnight rebuild

One evening's approval, "you're go", started an unattended overnight run. Four stages, no human in the loop until morning.

1 **Scrape**

An agent fan-out captured every live page and post as structured JSON (sections, copy, CTAs, links), an asset manifest, metadata, existing schema, and screenshots for QA. The preserved raw HTML later paid for itself repeatedly during remediation.

2 **Foundation**

Design tokens extracted from the live theme CSS, a component library of one module per pattern, and a SectionRenderer that maps typed content blocks to components.

3 **Assembly**

Agents wrote content files only. **No agent ever wrote a component.** 54 marketing pages, then 544 blog posts converted to MDX.

4 **QA**

Build green, every internal href resolved, metadata audited against the scrape (one cosmetic diff across 54 pages), screenshot comparison against the live originals, and a slug-parity check against the crawled URL inventory.

By morning, from a standing start.

703	100	100	1
STATIC PAGES GREEN	LIGHTHOUSE PERF	LIGHTHOUSE SEO	NIGHT

Lighthouse figures are sampled pages on a local build. Two daytime fidelity passes followed: the overnight build had styling drift, so the design system was rebuilt from the live theme CSS and all 53 pages re-verified side by side.

SCOPE NOTE, FOR HONESTY

The estimate said roughly 50 blog posts. The crawl found **544**. The fan-out architecture made a 10× scope error absorbable inside the same overnight window, which is a property worth designing for on purpose.

PHASE 2 · 5 TO 6 JULY

The SEO, GEO and AEO layer

A dedicated pre-launch plan, 46 findings ordered by severity, implemented in a day.

Preserving equity at the seam

Zero-redirect slug parity for 592 URLs, plus a 301 map for the 46 URLs that could not be preserved, so no equity died at cutover.

The full JSON-LD stack

Organization (foundingDate, founder, address, contactPoint, awards, and sameAs to the HubSpot directory, Clutch and LinkedIn), WebSite, Service on service pages, BreadcrumbList sitewide, FAQPage wherever an FAQ block exists, and BlogPosting plus Person on every post.

Answer-engine plumbing

llms.txt and llms-full.txt, a robots.txt that **explicitly allows GPTBot, ClaudeBot, PerplexityBot and Google-Extended**, and per-page question-form H2s matching the answer-engine query space. Every money page title rewritten too: primary query left-anchored, one intent per page, every description ending in a proof differentiator.

The blog template, rebuilt to an AEO spec

Extraction blocks

- Quick Answer box, 40 to 60 words, with at least one stat or named entity
- Key Takeaways
- Up to six FAQs emitting FAQPage schema **character for character**

Trust and structure

- Breadcrumb and schema, linked bylines with Person schema and LinkedIn sameAs
- Visible dateModified, reading time, auto-generated TOC with stable anchors
- A Sources block for citations, contextual CTA band, related posts

THE QUICK ANSWER BLOCK IS THE LOAD-BEARING ONE

It is the block AI search extracts. Every one of them is 40 to 60 words and is required to contain a statistic or a named entity, because an answer with nothing checkable in it is not worth citing.

PHASE 3 · 6 TO 7 JULY

The corpus rewrite

The centerpiece. Migrating 551 legacy posts unchanged would have relaunched the same content farm on faster hosting.

Pass 0: deterministic cleanup, no model involved

A scripted pass over all 544 posts, built from an empirical junk catalog: **468 broken image paths, 1,500   artifacts, 724 un-decoded HTML entities, 356 bold-lines-as-headings, 179 broken heading hierarchies**, absolute self-links, empty links and raw HTML wrappers.

Scripts fix mechanical rot. Models never touch what a regex can prove.

Pass 0b: FAQ extraction

168 posts already carried inline FAQ sections. A deterministic extractor moved them into frontmatter so the template renders them as accordions and emits FAQPage schema. **133 posts shipped that way.**

Data-driven triage, not vibes

Search Console over 16 months, GA4 and backlink data joined into a per-post disposition sheet: keep-and-spin, keep-as-is, merge, 301 or 410, with a target keyword for every surviving post.

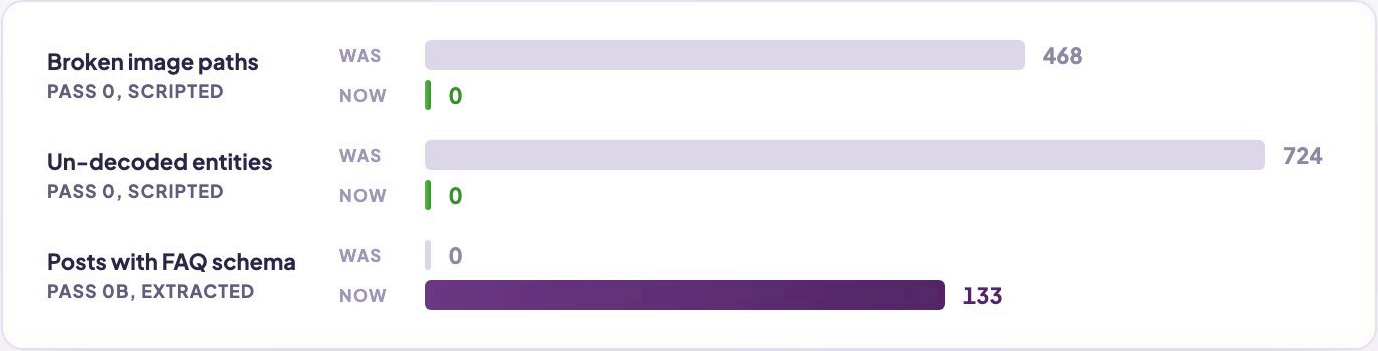

THE KILL CRITERION FOR THE REWRITE

A finished post must contain agency-specific substance that a generic version could not contain. **Pronoun-swapping is rejected as a fake spin.** That single rule is what separates a corpus rewrite from a mass paraphrase, and it is enforced by a different model than the one that wrote the draft.

THE REWRITE PIPELINE

Three passes, three different jobs

Per post, pipelined across 16 gated batches. The model that writes is never the model that verifies.

INPUTS, PER POST

- **Triage sheet** — Disposition + target keyword
- **Source library** — 2,065 vetted proof points
- **Link map** — Internal link graph
- **Nomenclature** — HubSpot naming rules
- **Canonical facts** — The stats of record

PASS 1 · OPUS — ICP re-angle

- Retell every generic marketing how-to from the agency-delivering-for-clients perspective, on the same URL
- Answer-first: every H2 answered in its first two sentences
- Statistics, quotations and cited sources woven in from the source library
- Writers select evidence. They never research or invent mid-rewrite

›

PASS 2 · SONNET — Distillation

- Quick Answer, 40 to 60 words, stat or entity mandatory
- Key Takeaways
- Up to six FAQs with standalone-extractable answers
- Image alt text
- Tags remapped to a rationalized 14-tag taxonomy, down from 33
- CTA selected from the library

›

PASS 3 · SONNET — Adversarial verify

- A fresh agent that never wrote the post
- Verifies facts against the canonical stats
- Traces every citation back to the source library
- Hunts fake spins
- Returns PASS, FIX or RE-ANGLE

RE-ANGLE returns the post to Pass 1 — PASS releases it into the wave

Wave gate between every batch

Build green, JSON-LD validation, internal links resolved, human spot-reads on the early waves. One commit per wave, pilot batch first.

16 GATED BATCHES

Batch 1 committed 21:09 on 6 July.
Batch 16 finished 16:24 on 7 July.

448	19.5	2,065
POSTS REWRITTEN	HOURS	PROOF POINTS

Includes two evidence-backfill passes that wove externally cited statistics into the early batches. Order-of-magnitude budget for the whole corpus rewrite: 6 to 9 million tokens, dominated by the Pass 1 rewrite.

PHASE 4 · 7 JULY, THE SAME EVENING

Media at corpus scale, at zero inference cost

Both pipelines ran on local models on Apple Silicon. The marginal cost of the second one is the same as the first: nothing.

551 featured images regenerated

- A local **FLUX.2** model (flux2-klein-4b via mflux) paints the background scene
- The per-post scene prompt is written by a small **local LLM**
- A compositor overlays brand geometry, category label and typeset title at 1200×628
- Prompts force the subject into the right third, because the brand overlay reveals the photo through a right-side diagonal window
- A color-rotation pass prevents runs of same-colored cards in the blog grid

459 posts narrated as audio

- Markdown stripped to speakable text
- Synthesized with **Kokoro TTS**, locally
- Encoded to 64kbps mono MP3, roughly 1MB per two minutes
- Uploaded to Cloudflare R2
- Served by a **zero-JavaScript** native audio player in the post sidebar
- 486 posts carry audio today; new posts get narration as they publish

Waves at 19:52 and 21:08. Audio manifest completed 23:59.

$0	$0	1,010
TO GENERATE	ONGOING	ASSETS IN ONE EVENING

NEITHER OF THESE IS DECORATION

The images fixed a corpus in which **28% of body images rendered broken**. The audio adds a modality that AI assistants and accessibility tools both consume, on content we had just spent 20 hours making worth consuming.

The reason to run these locally is not only the bill. A local pipeline has no rate limit and no per-call anxiety, so regenerating the whole corpus a second time after a design change costs an evening of wall-clock and nothing else. That changes what you are willing to attempt.

PHASE 5 · 8 JULY, THEN 30 JULY TO 6 AUGUST

Cutover, and the error purge

DNS cut over on 8 July with slug parity, the redirect map, 410s for deliberately retired content, and the sitemap submitted to Search Console the same day. Then the inherited debris.

The 511-URL "Not found" export went to zero remaining

408 URLs now answer 410 (deliberate removals and HubSpot platform bundles, matched by shape so the open-ended set stays covered), **95 became topical 301s**, 6 were live, and renamed tag archives got one-hop 301s.

> THE 410-VERSUS-301 EDITORIAL RULE, NOW MECHANIZED IN CODE
>
> A deliberately removed URL answers **410** so Google retires it. Only URLs with current traffic or links earn a **301**, and only to a genuine topical match. Redirecting dead content to an index page is a soft 404 and is never done. The old site had 31 of 58 redirect rules in the soft-404 pattern. The new config has zero redirects pointing at /blog.

What a full crawl of all 709 sitemap URLs turned up

- **25 posts carried the old page's scraped chrome inside the body**: title, author card and a "Latest Articles" module rendered twice. All repaired, one rebuilt from the preserved scrape.
- **26 unclickable CTA banner images**, which were pictures of buttons, removed.
- **23 dead outbound links** fixed or unwrapped, verified with browser-style GETs over IPv4 after HEAD-based scanning produced false positives.
- A broken og:image on two case studies, so every LinkedIn share had been rendering naked; a 45MB video auto-streaming on page load; and three buttons still routing through HubSpot's CTA redirect, including a "Download Free Ebook" button that delivered no ebook.

Every fix became a permanent gate

The content validator now runs in **npm test**, checks pages as well as posts, and asserts that its own route sets are non-empty. A validator bug had been reporting **948 phantom broken links**, which is exactly what let the 26 real ones hide. Each new gate was proven to fire against an injected violation before it was trusted.

Measurement rebuilt, 8 to 9 August

GTM host-gated, GA4 key events cleaned, custom dimensions added, and **94% of Total Blocking Time traced to third-party tags** (CallRail, Clutch, and two pixels HubSpot's tracking script was injecting). Removing them took homepage TBT from 1,630ms to roughly 130ms.

THE VERIFICATION CULTURE

Speed came from agents. Safety came from gates.

Six rules, every one of them learned the expensive way. They are the difference between a corpus rewrite and a corpus incident.

Waves with gates, never big bangs

Pilot batch first, human spot-reads, scripted validation between every wave, one commit per wave. A bad wave costs one wave.

The writer never verifies

Pass 3 is always a fresh agent that has never seen the draft being checked, and never wrote it.

Deterministic work never uses a model

And the inverse is just as strict: model output always faces a deterministic check.

Suspect the checker before the content

When a checker fails at scale, it is usually broken. A validator reported 948 phantom broken links, which is what hid the 26 real ones. Every derived set now asserts non-empty.

Prove that gates fire

By injecting a violation and watching it fail. A gate that has never failed is an unverified gate, not a clean codebase.

Screenshots verify, curl is ground truth

A vision-agent fan-out re-bills every screenshot in context on every tool call. One review loop burned 488M context tokens before being redesigned text-first at roughly 5% of the cost, for the same output.

THE BUDGET, STATED PLAINLY

Order of magnitude for the whole corpus rewrite: **6 to 9 million tokens**, dominated by the Opus rewrite pass. The scrape, media and audio pipelines ran on local models at zero marginal cost.

III

PART THREE

Results, one month post-cutover

All figures from Google Search Console and GA4, captured 6 to 10 August 2026. Every claim below is a screenshot in this document, and the ones that went the wrong way are here too.

INDEXATION AND CRAWL

A 15x expansion of the surface Google is willing to serve

From roughly 40 indexed pages to 620 in four weeks, on a domain Google had all but stopped crawling.

METRIC	PRE-MIGRATION TROUGH (MAY TO JUNE)	NOW (6 TO 10 AUGUST)
Indexed pages	roughly 30 to 40	**620**
Crawl rate	50 to 100 requests/day	**500 to 1,200 requests/day**
Not found (404), affected pages	951	**129 and draining**
Crawled, currently not indexed	roughly 1,150	**448**
Sitemap	31 URLs at its worst, against roughly 670 junk-heavy known URLs	**687 clean URLs, zero noindex, zero errors**

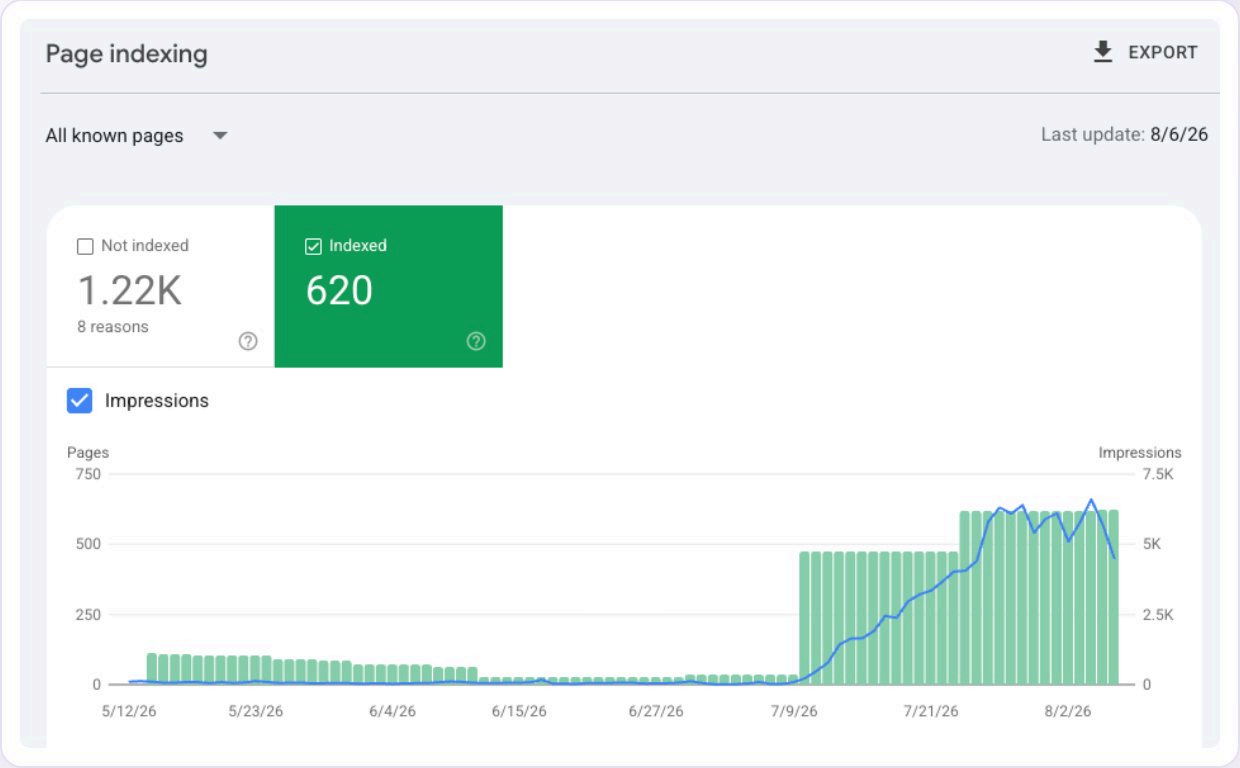

FIG 2 Page indexing. The index refills within days of the 8 July cutover. Source: Google Search Console.

THE INHERITED DEBRIS

The error backlog draining

Both curves below are the 410-versus-301 rule taking effect: deliberately removed URLs answer 410 and Google retires them, rather than 404-ing and being recrawled for months.

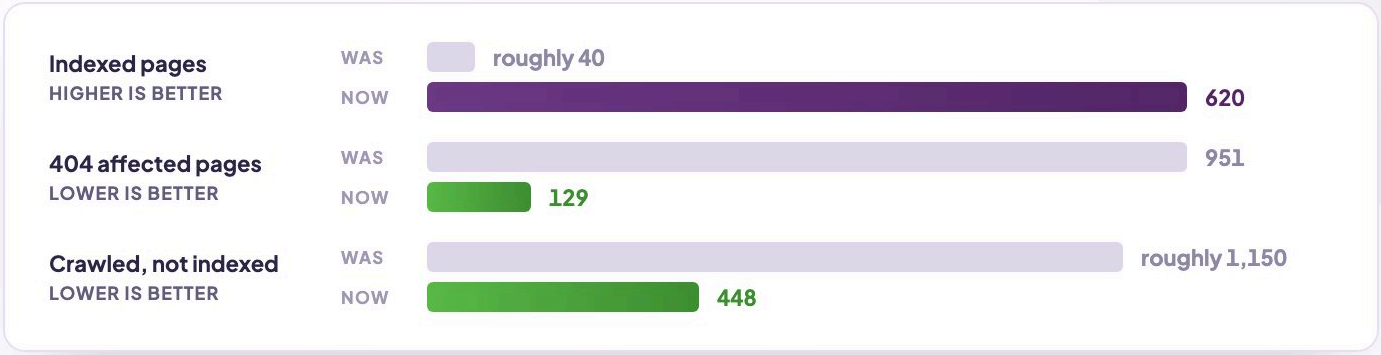

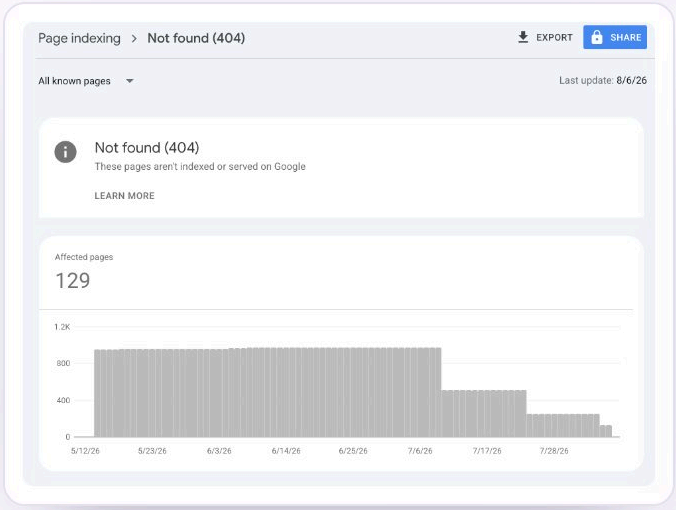

FIG 3 Not found (404): roughly 950 in June to 129 in early August, in two visible steps. Source: Google Search Console.

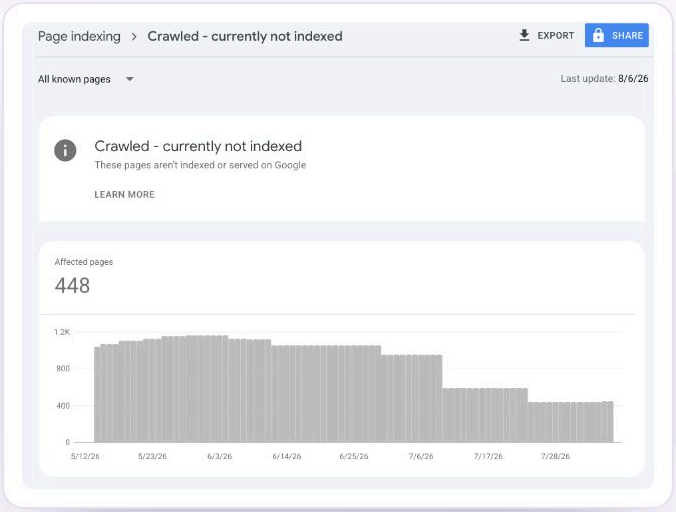

FIG 4 "Crawled, currently not indexed" halving from roughly 1,150 to 448. Source: Google Search Console.

WHY 410 AND NOT 404

A 404 tells Google the URL is missing, which it treats as possibly temporary, so it keeps recrawling for months and spends crawl budget you need elsewhere. A **410 is a statement that the resource is gone on purpose**, and it retires far faster. On a site with a 940-URL backlog and 50 crawls a day, that distinction is the whole recovery.

SEARCH PERFORMANCE

A flat line until 8 July, then a wall

+1,010%
CLICKS, 28 DAYS
111 clicks, www property

+5,335%
IMPRESSIONS, 28 DAYS
119K impressions

25,720
WEEKLY IMPRESSIONS
Week 3, from 976 in week 1

5–7K
IMPRESSIONS PER DAY
Current run rate

Weekly impressions went **976, then 12,250, then 25,720** across the first three weeks after cutover, and now run at roughly 5,000 to 7,000 a day. The three-month performance chart below is the clearest single artifact in this document: nothing, then the cutover date, then a wall.

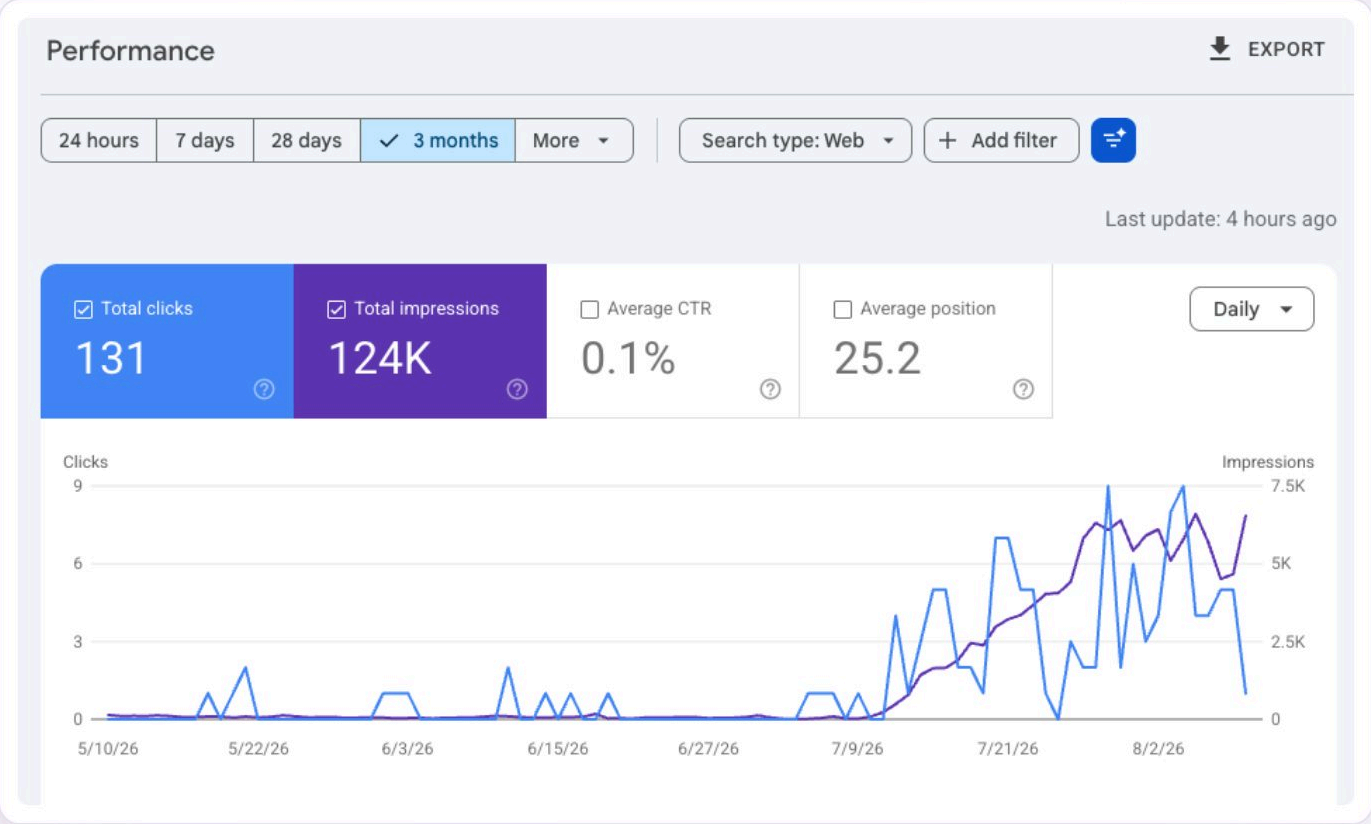

FIG 5 Performance, three months: 131 clicks and 124K impressions, both series flat until 8 July and climbing steeply after. Source: Google Search Console.

All figures are the www URL-prefix property. Property choice matters and is stated on every number in this document: the apex domain property kept serving the old index through the trough, so the sharp recovery line is www-specific.

THE CAVEAT, STATED PLAINLY

Average position got worse. That is the good news.

Average position moved from roughly 15 to roughly 25 at cutover. This is a composition effect, not a rankings drop, and telling the difference is the single most valuable habit in this whole document.

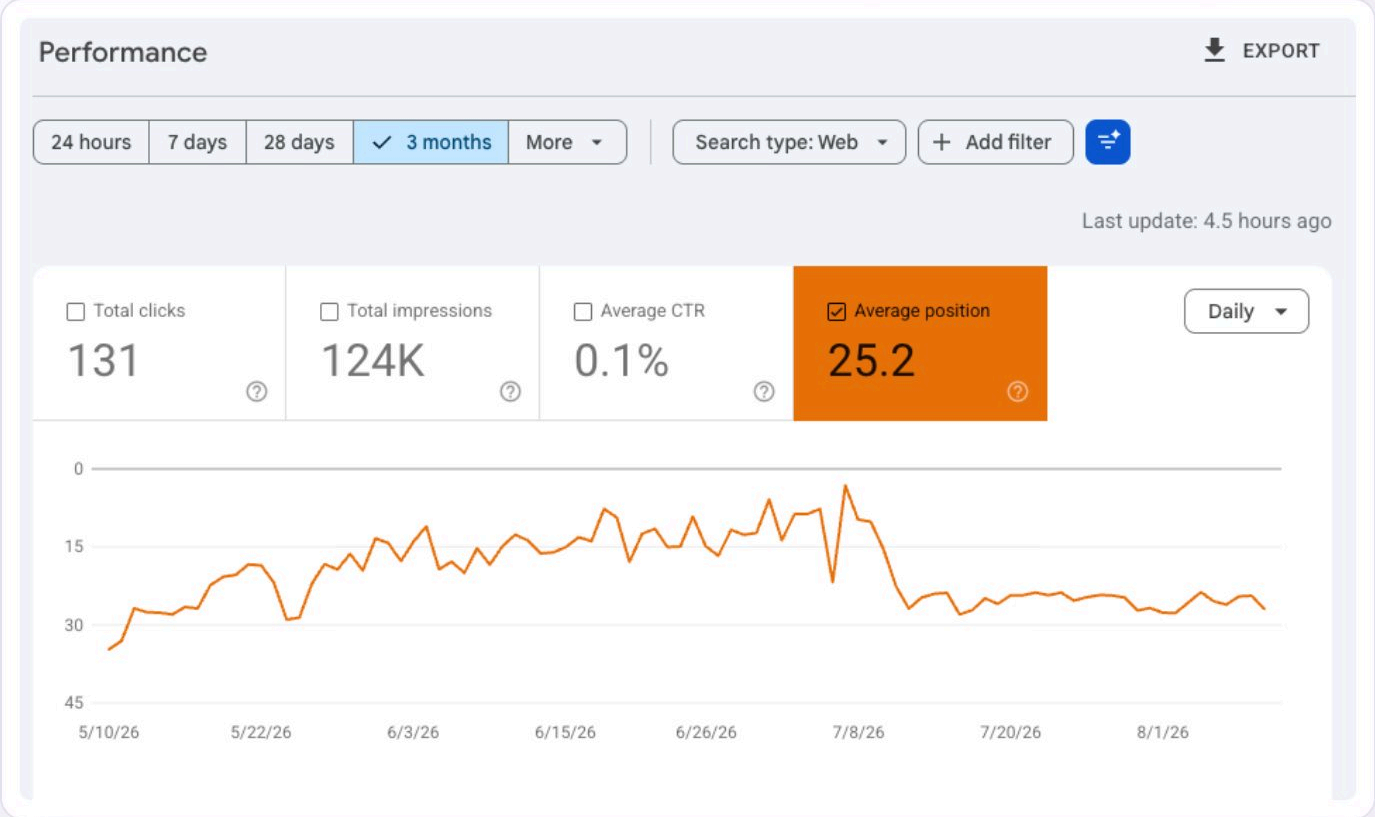

FIG 6 Performance with average position selected: the average moves from around 15 to around 25 exactly as hundreds of new pages enter the index. Source: Google Search Console.

Before cutover, the average was computed over a tiny surviving footprint, a few dozen pages that had held on. At cutover, hundreds of pages entered the index at deep initial positions. Every one of them dragged the average down while **adding** impressions. The metric got worse because the business got better.

THE SAME ARITHMETIC, TWICE, IN OPPOSITE DIRECTIONS

In 2025 this site's dashboard showed "top-10 keyword counts at a multi-year high" while clicks fell every month. That was a false comfort. In 2026 it shows average position worsening while impressions grew 50×. That is a false alarm. **Averages lie at both ends of a migration.** Watch clicks and conversions per query class instead.

AI VISIBILITY

From absent to present inside Google's AI answers

The Generative AI features report in Search Console was effectively zero before 12 July. By early August it peaked near 240 impressions a day, 2,060 in the window. The ramp starts almost exactly at cutover.

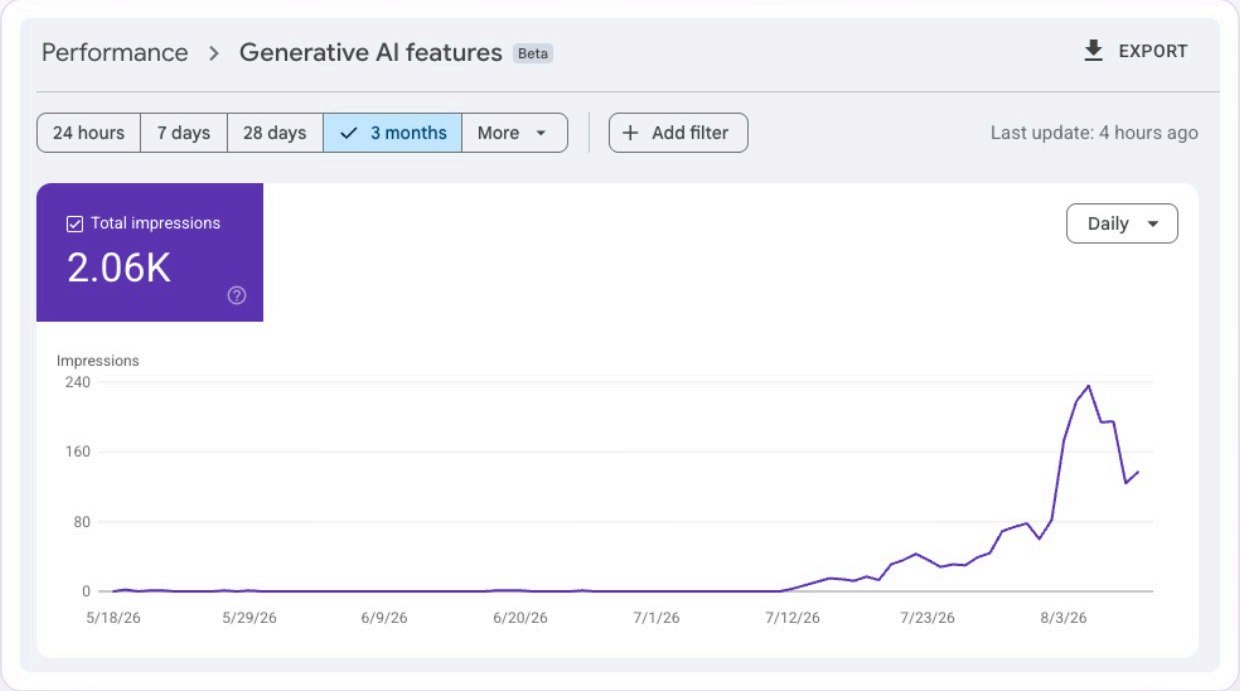

FIG 7 Generative AI features: 2.06K total impressions, a flat zero line until mid-July, then a steep ramp peaking near 240 impressions per day in early August. Source: Google Search Console.

AI-assistant referrals nearly tripled

GA4 SESSIONS FROM AI ASSISTANTS	JUNE	JULY
All assistants (ChatGPT, Claude, Gemini, Perplexity, Copilot)	18	53
claude.ai specifically	roughly 2	12

For context and against our own interest: the old site saw roughly 100 AI-referral sessions a month in mid-2025, earned by a 551-post consumer corpus at 5x today's total traffic. The new curve is being earned by ICP pages a buyer can actually convert on, which is a different and more useful number.

AND THE ANSWERS THEMSELVES

Three engines, one positioning, all citing the site's own pages

Impressions prove presence. These prove placement. All captured 10 August 2026 on buying-intent queries.

ChatGPT *"What's the best white label HubSpot agency for partner agencies?"*

"My top pick right now is Meticulosity." Ranked Best overall.

Cites the exact proof stats the rewrite wove into the site: Diamond Solutions Partner at a 5.0 marketplace rating, 12+ years a HubSpot partner, 17+ years as an agency, and working behind the scenes for partner agencies rather than competing for the end client.

Perplexity *"Who provides white label HubSpot support for agencies?"*

Places Meticulosity first in its provider table, across ten sources, with citations to meticulosity.com carried inline.

Claude *"Who is the best white-label HubSpot agency?"*

Leads its list with Meticulosity: "the most white-label-native of the bunch, it's their whole business, not a side line."

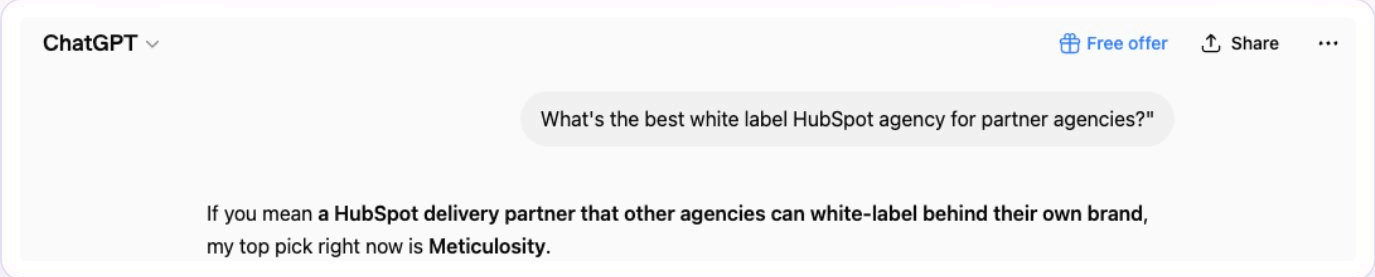

FIG 8 The top of the ChatGPT answer, uncropped from the first pixel: the query as asked, and the verdict. The full answer, and the Perplexity and Claude captures, are held with the source data. **Source: ChatGPT, captured 10 August 2026.**

WHY THIS IS THE PAYOFF, NOT A VANITY METRIC

Nobody optimizes for "be named first by ChatGPT". What was optimized for was answer-first copy, FAQPage schema, cited statistics and open AI crawlers. Placement is what that produces when it works, and the citations point at the site's own pages, not a directory listing.

THE NUMBER THAT PAYS FOR EVERYTHING

Leads, and page performance

We pulled the complete form-submission history from **both** HubSpot portals, the legacy portal that served the site until the early-2026 portal migration and the current one, summed across the transition, and filtered to the site's three lead forms.

PERIOD	SITE LEAD-FORM FILLS PER MONTH
2025 monthly range	12 to 33, averaging roughly 22
Q4 2025	28 to 33
April to June 2026, the collapse trough	23, then 14, then 12
July 2026, first month on the new site	**22, nearly double the June trough**

The honest read: one month in, lead-form volume has recovered from half-rate back to the long-term average, on impressions that are still compounding. Clicks are the leading indicator and are still climbing through the 4 to 12 week content window, so the leads line is expected to follow. That is what the 60 and 90-day updates are for.

Measurement note, because this is the money claim: GA4 event months before August 2026 undercount, since mount-fired events were silently dropped until the 9 August fix, and its conversion tag was dead from April to June. HubSpot form counts are therefore the instrument used here, summed across both portals through the migration seam. GA4 now also tracks phone-call conversions, new as a measured event, 10 in the trailing 60 days.

Performance: the same-week A/B run

SAME TEST RIG, 4 TO 6 JULY 2026	HUBSPOT (PRODUCTION)	NEXT.JS (NEW BUILD)
Homepage GTmetrix Performance	70%	**99%**
Homepage Structure	85%	**99%**
Homepage LCP / TBT / CLS	1.4s / 568ms / 0.08	**726ms / 41ms / 0**
Money page Performance	69%	**98%**
Money page LCP / TBT	1.6s / 467ms	**1.0s / 90ms**
Money page TTFB / fully loaded	459ms / 3.6s	**95ms / 1.9s**

Since launch, third-party tag removal cut homepage Total Blocking Time a further 92%, from 1,630ms to roughly 130ms. CLS is 0 everywhere and Lighthouse SEO scores 100/100 on every page tested. Mobile lab runs on the live site vary between 52 and 97 across single runs, a documented variance under simulated throttling; field data is the metric that actually ranks and is being enabled.

IV

PART FOUR

What we would tell another team

Seven things, in the order they would have saved us the most time. None of them is “pick this framework”.

SEVEN LESSONS

What we would tell another team

1 **The migration is the smallest part**

Rehosting the same corpus buys a faster version of the same decline. Budget the rebuild of the content, not just the site.

2 **Traffic that cannot become a customer is a liability**

The old site's biggest page ever, 5,164 clicks in its final 16 months, never produced a lead and was killed on purpose. Deliberate removals ship as 410s.

3 **Make content machine-operable**

One typed file per page and per post is what made 448 rewrites, 551 images and 459 narrations scriptable in hours. A closed CMS makes every one of those a project.

4 **AI does the work; gates make it shippable**

Multi-pass pipelines with a separate adversarial verifier, wave gates, injected-violation tests for every new check, and humans spot-reading the early waves.

5 **Answer-first is the new on-page**

Quick Answers, question-form H2s, FAQPage schema, cited statistics (the Princeton GEO trio: statistics, quotations, sources), entity-complete Organization schema, llms.txt, and open AI crawlers. Google's Generative AI impressions and assistant referrals both responded within weeks.

6 **Averages lie at both ends of a migration**

Keyword counts said the old site was fine while clicks halved. Average position says the new site got worse while impressions grew 50×. Judge on clicks and conversions per query class.

7 **Instrument before you need it**

The Search Console domain property, clean GA4 key events and a repeatable crawl and audit harness are what turn "it feels better" into this document.

WHAT HAPPENS NEXT

We publish the follow-ups, including what goes wrong

The effect window for content work at this scale is 4 to 12 weeks, which means everything in Part III is the **early** return. Waves of rewritten posts are still being recrawled and the "crawled, currently not indexed" pool is still draining.

We will publish updated figures at the 60-day and 90-day marks, mid-September and mid-October 2026, in this document. That includes the metrics that got worse or went nowhere. If the thesis is right, the compounding shows up there. If it is wrong, that is worth publishing too.

Content quality, verified independently

An independent MIT-licensed audit toolkit (claude-seo, not used to build any of the remediation) scores the money pages **84 to 92 on content quality with zero filler-phrase and zero AI-pattern flags**, agent-UX 100/100 for real buttons, real anchors and semantic landmarks, no soft-404s, and no AI crawler blocked. A fresh audit on 10 August 2026 shows zero critical and zero warning regressions against the 1 August baseline.

84–92	0	100	0
CONTENT QUALITY SCORE	AI-PATTERN FLAGS	AGENT-UX SCORE	DRIFT REGRESSIONS
Money pages, independent toolkit	On a corpus rewritten by AI	Real buttons, anchors, landmarks	Against the 1 August baseline

THE CLAIM WE ARE MOST WILLING TO BE HELD TO

Zero AI-pattern flags on a corpus that was rewritten by AI is the finding we would most like other teams to test us on. It is not a defense of the tooling. It is evidence that the **gates**, not the models, are what determine whether machine-written content reads like machine-written content.

Next updates to this document.

60d	90d	All
MID-SEPTEMBER 2026	MID-OCTOBER 2026	METRICS, INCLUDING THE BAD ONES

POSTSCRIPT

This is not an anti-HubSpot paper

We are a Diamond HubSpot partner. The CRM, forms, meetings and marketing automation all stayed on HubSpot by choice, and they are staying. We left one product, Content Hub, for one reason: our content had to become something AI agents could operate on at corpus scale, and in mid-2026 there was no supported way to do that inside it.

Be concrete about what that gap costs. Here is this exact project, priced in time against Content Hub as it exists today. Scoping this kind of work is our day job, so these are the bands we would quote a client.

WORKSTREAM	WHAT IT TOOK US	ON CONTENT HUB TODAY
Rebuild 54 marketing pages, SEO and AEO optimized	**One overnight run**	Pages are drag-and-drop module trees, and the Pages API cannot safely rewrite layout_sections content, so this is hand-work in the editor: **2 to 6 hours per page, 6 to 10 weeks elapsed**
Rewrite 448 blog posts: ICP re-angle, Quick Answers, FAQs, citations	**~20 hours, 16 batches**	The Blogs API can write body HTML and meta, so a pipeline is *possible*. But Quick Answers, takeaways, FAQ schema and sources have no home without custom HubL modules plus HubDB wiring per post. With a custom pipeline: **4 to 8 weeks**. By hand: **900+ hours**
Template layer: TOC, answer blocks, author E-E-A-T, full JSON-LD stack	**Days, in code, with tests**	**2 to 4 weeks** of HubL and module development, with no local test harness for the schema output
551 branded featured images	**One evening, $0**	Generation is external either way; attaching via API is **days**, via the UI **about a week**
459 audio narrations plus the player	**One evening, $0**	Hosting is external either way; the player is another custom module plus per-post wiring
Error purge: true 410s, bundle URL spaces, wave-gated rollouts, validators in CI	**Days, permanent gates**	**Largely not possible.** URL mappings do 301 and 302 only, the platform's own JS/CSS and AMP URL spaces are not yours to control, and there is no corpus-wide dry-run, diff or rollback

Add it up and the honest answer is that **what took six days would have been one to two quarters with a dedicated team**, and several workstreams (status-code control, CI validation, atomic corpus-wide waves) have no path at any price. Not because HubSpot is bad software, but because Content Hub was designed for humans editing one page at a time, and the leverage in 2026 is agents editing five hundred pages at a time behind verification gates.

THE OFFER

The day HubSpot ships first-class agentic access to Content Hub, we would go back in a heartbeat.

The wishlist is short and specific.

1. A **bulk content API** that treats pages and posts as structured, typed content rather than opaque module trees, with dry-run diffs and transactional, revertible batches.
2. **Custom structured fields on blog posts**, the equivalent of frontmatter, so answer blocks, FAQs and citations are data rather than template hacks.
3. **Real URL-space control**: 410s, and ownership of every URL the platform emits under your domain.
4. **An MCP server for Content Hub**, so any agent, not just a custom pipeline, can read, propose and stage content changes behind human approval.
5. **Webhook or CI hooks on publish**, so validators can gate content the way tests gate code.

WHAT THAT WOULD BUY

Give us that and everything in this paper becomes a **Content Hub playbook**: same pipelines, same gates, same six days, with HubSpot's CRM attribution native instead of stitched. We would be the first partner to run it, and we say that as the agency that just left.

APPENDIX A

Technical inventory: what shipped, specifically

Stack and content model

Next.js App Router, TypeScript, Tailwind 4. Static generation for all roughly 755 built pages. Content as typed JSON blocks (105 marketing pages) plus an MDX frontmatter contract (557 posts).

Structured data

Organization (founder, foundingDate, address, contactPoint, awards, sameAs), WebSite, Service, BreadcrumbList on 639 pages, FAQPage on 484 pages plus 133 extracted post FAQs, BlogPosting with dateModified, wordCount, articleSection and keywords, and Person with @id and sameAs.

The AEO template, per post

Quick Answer (40 to 60 words, stat or entity mandatory), Key Takeaways, TOC with stable anchors, up to six FAQs matched to schema character for character, Sources list, linked author bios, visible Updated date, reading time, audio player.

Answer-engine surface

llms.txt and llms-full.txt. robots.txt explicitly allowing GPTBot, ClaudeBot, PerplexityBot and Google-Extended. Sitemap with lastmod on all 687 URLs.

Redirect and 410 machinery

A 46-URL launch redirect map, a GONE list plus HubSpot bundle-path 410s matched by shape, one-hop 301s for renamed tag archives, and zero redirects pointing at index pages.

Rewrite pipeline

Pass 0 scripted cleanup against a 12-class junk catalog, Pass 0b FAQ extraction, Pass 1 Opus ICP re-angle, Pass 2 Sonnet distillation, Pass 3 Sonnet adversarial verify. A 2,065-item source library, a 14-tag taxonomy, and per-post keyword targets from the triage data.

Media

Local FLUX.2 image generation plus a brand compositor at 1200×628 progressive JPEG. Kokoro TTS narration to Cloudflare R2 with a manifest-driven player.

Validation in CI

Content validator (route parity, image and link integrity, the frontmatter contract, branded-image detection, gone-list consistency), a metadata audit against the original scrape, a full-crawl link checker, and JSON-LD validation.

Measurement

GA4 and GTM, host-gated, with event-scoped custom dimensions. Search Console URL-prefix and domain properties. A weekly in-house SEO monitor and a monthly audit battery: the in-house toolkit, an independent claude-seo run, and drift comparison against a baseline.

APPENDIX B

Timeline

Every date below comes from the site repository's git history.

2 July, overnight — Scrape and full rebuild
703 static pages building green by morning, Lighthouse 100 on performance and SEO on sampled pages.

2 to 3 July — Visual fidelity passes
The design system rebuilt from the live theme CSS; all 53 pages re-verified side by side against the originals.

5 to 6 July — SEO, GEO and AEO plan implemented
46 severity-ordered findings shipped; the blog template rebuilt to the AEO spec.

6 July, 17:35 — Pass 0 cleanup of all 544 posts
Deterministic junk removal across a 12-class catalog. No model involved.

6 July 21:09 to 7 July 16:24 — 448-post ICP rewrite
16 gated batches, three passes per post, adversarial verification on every one.

7 July, evening — 551 featured images and 459 audio narrations
Local FLUX.2 and Kokoro TTS on Apple Silicon. Zero inference cost.

8 July — DNS cutover, sitemap submitted
Slug parity for 592 URLs, a 301 map for 46, and 410s for deliberately retired content.

30 to 31 July — Post-launch audit
Link, chrome, CTA and outbound purge across all 709 sitemap URLs. Validator hardened.

6 August — Search Console 404 backlog cleared
The 511-URL export went to zero remaining.

8 to 9 August — Measurement stack rebuilt
Third-party tags purged; homepage Total Blocking Time down 92%.

10 August — This snapshot
620 indexed, clicks +1,010%, impressions +5,335%.

meticulosity

THE HUBSPOT AGENCY FOR AGENCIES

We do this for other agencies' clients too.

Meticulosity is a white-label HubSpot partner working invisibly under your brand: migrations, portal audits, onboarding, automation and delivery. Everything in this document, we ran on ourselves first.

70+	11,800+	17+	Diamond
AGENCY PARTNERS	PROJECTS DELIVERED	YEARS IN BUSINESS	HUBSPOT PARTNER

meticulosity.com ›