KOOKIES

GUIDE

The AI-built Site Rescue Guide

6 CHAPTERS

01 What AI builders get right and what they leave o

02 The 20 minute check: can Google and AI assistant

03 SEO basics most AI-built sites miss

04 Speed and phones

05 Fix in place, render on the server, or rebuild

06 Moving to the new site without losing traffic

SEO AND AEO

The AI-built Site Rescue Guide

Start by checking whether your words are in the first HTML and what Google's URL Inspection tool renders, because many AI crawlers do not run JavaScript. Then fix titles, the sitemap, images and scripts in place, switch to server rendering if your tool offers it, or rebuild in Webflow or Astro when rendering, editing and speed are all problems. If URLs change, a 301 for each one keeps what the site has earned.

Who it is for. Founders and marketers whose site was built with Lovable, Bolt, v0, Framer AI or a similar AI tool, and who want it found, fast and easy to edit.

01 What AI builders get right and what they leave out

02 The 20 minute check: can Google and AI assistants read it?

03 SEO basics most AI-built sites miss

04 Speed and phones

05 Fix in place, render on the server, or rebuild

06 Moving to the new site without losing traffic

Read online at www.kookies.io/guides/ai-site-rescue. Updated 2026-11-16.

CHAPTER 1 OF 6

What AI builders get right and what they leave out

01

What the tools do well, the gaps a marketing site still has, and the three kinds of AI-built site.

The short answer

AI builders such as Lovable, Bolt, v0 and Framer AI are very good at the first version: a layout, a look and working pages in an afternoon. What they tend to leave out is the work a marketing site needs after that. Pages that search engines and AI assistants can read in the first HTML, a title and description per page, a sitemap, redirects, speed on a mid-range phone, forms that really deliver, and a way for your team to edit copy without prompting again.

None of that means the tool was a mistake. It means the site is a prototype that went live. This guide helps you find out which gaps your site has, fix the ones you can fix in place, and decide when a rebuild is the cheaper path. If you only want a quick test first, our post on checking if Google can read a Lovable, Bolt or v0 site covers the basics in seven minutes. This guide goes further: what each result means, what to fix first and how to move without losing traffic.

What the tools do well

Give credit where it is due. For a founder or a small team, these tools solve real problems:

- **Speed to a first version.** You can show an idea to buyers or investors the same day.
- **A design you can react to.** It is easier to say "more space here, a darker hero" to a real page than to a brief.
- **Working parts.** Forms, pricing toggles and small interactive pieces often work out of the box.
- **Code you can take with you.** Most of these tools let you export the code or sync it to GitHub, so you are not starting from nothing.

Keep all of that. A good rescue keeps the design you like and the copy that works. What changes is how the pages are built and served.

Where the gaps usually are

The gaps come from the same place: these tools are built to make an app that works when a person opens it in a browser. A marketing site has a second audience, the crawlers that decide whether anyone finds it. And a third, the team that edits it every week.

AREA	WHAT OFTEN SHIPS	WHAT A MARKETING SITE NEEDS
Rendering	An almost empty HTML file, with the text drawn by JavaScript in the browser	The words, links and headings in the first HTML
Titles and descriptions	One title and description for the whole site, or a generic default	A title and description written for each page
Discovery	No sitemap, or one that lists a preview address	A sitemap of real URLs on your real domain
Links	Buttons that change the page with a click handler	Real `<a href>` links between pages
Speed	Fast on a laptop on office wifi	Fast on a mid-range phone on mobile data
Editing	A new prompt for every copy change	A CMS or editor where a marketer changes words and publishes
Moving	No plan for old URLs	A 301 for every address that changes

Not every site has every gap. Some tools now render on the server by default. That is why the next chapter is a test, not a verdict.

Three kinds of AI-built site

Before you fix anything, name what you have. The right path depends on it.

1. **A marketing site.** Home, product or service pages, pricing, about, a blog, a contact form. No logins and no user data. This is the most common case and the one this guide is mostly about.
2. **An app with a marketing front.** A product people sign in to, plus a few public pages that sell it. The public pages need everything in this guide. The app behind the login does not need to rank.
3. **An app.** Logins, accounts, dashboards, saved data, payments. Search matters little here, and the work is about security, data and reliability.

For the first kind, we usually rebuild the public site in Webflow or Astro: see our AI-built site rebuild service. For the third kind, a marketing rebuild is the wrong tool. Talk to us about MVP development instead. For the second kind, split them: the marketing pages on a site built for search, the app on its own subdomain or path.

> DECISION RULE
>
> If a page has a login, saved data or payments, it is app work. If a page exists to be found and to convert a visitor, it is marketing site work. Plan them apart, even if they share a design.

Why the copy changes get harder

There is one gap that no crawler test will show. On many AI-built sites, the copy lives inside the code of each component. Changing a headline means a new prompt, and the tool rewrites part of the code to do it. Each change is small. Over months, the code drifts: spacing changes on pages nobody touched, a section loses its mobile layout, a fixed bug comes back.

That is fine for a prototype. For a site your team updates every week, it becomes slow and risky. A short test tells you where you stand:

- Can your marketing lead change a price on the pricing page without help?
- Can they add a blog post with a title, description and social image?
- Can they add a new landing page that looks like the rest of the site?
- After each change, did anything else on the site move?

If the answer to any of the first three is no, put editing on your list next to SEO. It often decides between fixing in place and rebuilding, which chapter 5 covers.

What this guide will and will not do

This guide gives you tests you can run yourself, the meaning of each result, and decision rules for what to do next. Every vendor fact links to the vendor's own page, because these tools change often. Check the page again before you act on it.

It will not promise rankings or AI citations. Nobody can. Google says in its JavaScript SEO guide that it processes JavaScript pages in three phases (crawling, rendering and indexing), and it does not promise when or whether any page is indexed. What you can do is remove the reasons a page cannot be read, and then measure.

Work through it in order. Chapter 2 is a 20 minute check. Chapters 3 and 4 are the SEO and speed basics that most AI-built sites miss. Chapter 5 is the decision: fix in place, render on the server, or rebuild. Chapter 6 is how to move without losing what the site has earned.

TAKEAWAYS

- AI builders are great at a first version. The gaps are in what crawlers, phones and editors need.
- Name what you have first: a marketing site, an app with a marketing front, or an app.
- Pages with logins and data are app work. Pages meant to be found are marketing site work.
- If your team cannot change copy without a new prompt, editing belongs on the list next to SEO.
- Nobody can promise rankings. You can remove the reasons a page cannot be read, then measure.

02 The 20 minute check: can Google and AI assistants read it?

Four tests in 20 minutes, what Lovable and v0 render today, and how to read the results.

The short answer

Pick three pages: the home page, your most important product or service page, and one blog post. For each one, check four things. Is your hero sentence in the page source? Does Google's URL Inspection tool show the full page in its rendered HTML? Is the title different from the other pages? Are the links real links? That takes about 20 minutes and tells you whether your problem is rendering, basic SEO, or neither.

The reason this matters: Google can run JavaScript, but in a later step. Many AI crawlers do not run it at all. A study by Vercel and MERJ published in December 2024 found that none of the major AI crawlers it looked at rendered JavaScript, including OpenAI's and Anthropic's. If a sentence is not in the first HTML, assume many AI assistants never see it.

Minutes 0 to 5: view source

Open each page in a browser, right click and choose View Page Source, or type `view-source:` before the URL. Do not use Inspect. Inspect shows the page after scripts have run, so it always looks fine.

Press Ctrl+F (Cmd+F on a Mac) and search for:

- ✓ A few words from your hero sentence
- ✓ A price or a number from the page
- ✓ One FAQ answer, if the page has FAQs
- ✓ The name of an inner page you link to, such as Pricing

Write down found or not found for each one. If the source holds a `<div id="root"></div>` and a script tag, and none of your words, the page is drawn in the browser. If the hero is there but the FAQ answers are not, part of the page loads later. Both are common. Our free AI crawler checker runs a similar test from outside your browser.

Minutes 5 to 12: what Google sees

Your own browser is not Google. To see what Google gets, use Search Console. If you have not set it up yet, verify your domain first. It is free and you need it for chapter 6 anyway.

1. Paste the page URL into the URL Inspection tool.
2. Run a live test, then choose View tested page.

3. Look at the screenshot, the HTML and the JavaScript console messages.
4. Search the HTML panel for the same four phrases as before.

Google's help page says the live test shows a screenshot, the HTML, response headers, console output and the page resources that loaded. Live tests have a daily limit per property, so test your three pages, not fifty.

No Search Console yet? Google's JavaScript SEO guide names the Rich Results Test as the other way to see rendered HTML. Paste the URL, wait for the test, and open the HTML view.

What Lovable and v0 render today

The answer depends on the tool and on when the project was made. These are the vendors' own words, checked in November 2026:

TOOL	WHAT THE DOCS SAY	WHAT IT MEANS FOR YOUR TEST
Lovable, new apps	Lovable's SEO page says apps created from 13 May 2026 (22 June 2026 on Enterprise plans) use TanStack Start with server-side rendering, so every request gets full HTML.	View source should show your words. If it does not, part of the page loads its data later.
Lovable, older apps	Older React and Vite apps get pre-rendering on deployed public URLs, served only to verified crawlers: search engines, AI search crawlers, social preview bots and verified SEO tools. People, unverified bots and link checkers get the single-page app.	Your view source may look empty while Google gets the full page. Trust URL Inspection over your browser here.
v0	v0's FAQ says it uses Next.js, React, TypeScript, Tailwind CSS and shadcn/ui, and publishes to Vercel.	Next.js can render on the server, but parts that fetch data in the browser can still be missing. Run the test.
Bolt and others	Depends on the framework the project uses.	Run the same test. Do not assume.

Lovable also says older projects can be upgraded to TanStack Start, that the published site does not change until you publish again, and that normal usage costs apply. If you exported your code and host it elsewhere, test again: Lovable describes pre-rendering for the URLs it serves.

Minutes 12 to 16: titles and links

Still in view source, check the head of each page.

- ✓ **Title.** Search for `<title>`. Is it different on each of your three pages? Google asks for distinct text that describes the content of the page, not repeated boilerplate.
- ✓ **Description.** Search for `name="description"`. Same check. Lovable's docs say a new project's starting metadata is generic.
- ✓ **Canonical.** Search for `rel="canonical"`. If it exists, it should point to the page's own address on your real domain.
- ✓ **Links.** Search for `href=`. Your menu items should be real links. Google's guide says it can only discover links that are `<a>` elements with an `href`. A button with a click handler is not a link to Google.

One more trap from the same guide: if your page has a `noindex` tag in the first HTML, Google may skip rendering, so removing it with JavaScript may not work. Search the source for `noindex` too.

Minutes 16 to 20: sitemap and robots

Open `yoursite.com/robots.txt` and `yoursite.com/sitemap.xml`.

- ✓ robots.txt loads and does not block the whole site with `Disallow: /`.
- ✓ It does not block the script and style files the page needs. Google's robots.txt guide says not to block resources that make the page harder to understand.
- ✓ The sitemap exists and lists your real domain, not a preview or staging address.
- ✓ Every page you want found is in it, and nothing you want hidden is.

Lovable says it creates a robots.txt with every new project, but not always a sitemap. Its SEO review checks for a missing or wrong `Sitemap:` line, relative URLs and URLs pointing at another host.

How to read your results

You now have a small grid: three pages, four checks each. Read it like this.

WHAT YOU FOUND	WHAT IT MEANS	WHERE TO GO
Words in view source, titles unique, links real	Rendering is fine. Your gaps are elsewhere.	Chapters 3 and 4
Empty source, but URL Inspection shows the full page	Google can read it. Many AI assistants may not, unless your host serves them a rendered copy.	Chapter 5, render on the server
Empty source and an empty or broken URL Inspection screenshot	Nobody can read it reliably.	Chapter 5, the full table
Same title on every page	Search results cannot tell your pages apart.	Chapter 3
Menu items are buttons, not links	Google may not find your inner pages at all.	Chapter 5

DECISION RULE

If a sentence matters to a buyer, it should be in the first HTML. Do not rely on a host serving crawlers a different copy, because you cannot see what every crawler gets.

TAKEAWAYS

- Test three page types, not just the home page. They are often built in different ways.
- View source shows the first HTML. Inspect shows the page after scripts, so it always looks fine.
- URL Inspection or the Rich Results Test shows what Google rendered. Trust them over your browser.
- Older Lovable apps serve a rendered copy to verified crawlers only. Your own view may differ.
- Many AI crawlers do not run JavaScript, so keep every sentence that matters in the first HTML.

CHAPTER 3 OF 6

SEO basics most AI-built sites miss

03

Seven basics, what good looks like for each, and a fix order.

The short answer

Seven basics cover most of the gap: a title per page, a description per page, one clear H1, a sitemap of real URLs, a robots.txt that lets crawlers in, a canonical URL on each page, and schema plus a social image where they fit. Each one must be in the first HTML, not added by a script after load. None of it is advanced. It is just easy to skip when a tool builds the page from a prompt.

Fix them in that order. Titles and descriptions change what people see in search results today. The rest helps search engines find and understand your pages.

1. A title for every page

Google uses the title element as a main source for the title link in search results. Its title link guidance asks for distinct text that describes each page, warns against repeated boilerplate, and says to brand titles concisely.

PAGE	WEAK	BETTER
Home	Kookies	Kookies: cookie boxes for teams and events
Pricing	Kookies	Pricing: cookie boxes from 12 to 500 | Kookies
Blog post	Kookies | Blog	How to plan treats for a 200 person offsite | Kookies

The examples use a made-up cookie shop. Write yours the same way: what the page is, then the brand. If you rebuild in Webflow, its help page on SEO titles and meta descriptions says SEO settings apply only to the page they are set on, and suggests keeping titles under 60 characters. Our free meta tag tool writes the tags and previews them.

2. A description for every page

The meta description does not change rankings much, but it can change what people read under your title. Google's snippet guidance says it mostly builds snippets from the page content, uses the description when it describes the page better, and that identical or similar descriptions on every page are not helpful.

- ✓ One to two plain sentences that say what the page offers and who it is for.
- ✓ Different on every page. Site-wide text suits the home page only.
- ✓ True to the page. A promise the page does not keep loses the click later.

3. One H1 and headings in order

AI tools often pick heading tags for their size, not their meaning. You end up with three H1s on one page, or an H4 straight after the H1 because it looked right. Screen readers and crawlers use headings to understand the page.

- ✓ One H1 per page that says what the page is.
- ✓ H2s for the main sections, H3s inside them. No skipped levels.
- ✓ Size is a style choice. Change the look in CSS, not the tag.
- ✓ Headings as text, not inside images.

Answer-first headings help AI answer tools too. Our AEO Readiness Guide covers how to write them.

4. A sitemap and robots.txt

Google's sitemap overview says you may not need a sitemap if your site has about 500 pages or fewer and every important page is linked from your home page. It also says a sitemap does not guarantee that every listed page is crawled and indexed. Add one anyway: it costs nothing, and it is how you tell Google about new URLs after a move.

- ✓ Lists only real, final URLs on your real domain, with the same form as your canonical URLs (trailing slash or not).
- ✓ Updates itself when you publish. Webflow can auto-generate one and update it on every publish. On Astro, the sitemap integration lists statically built pages; it cannot list dynamic routes rendered on demand, so add those yourself.
- ✓ Named in robots.txt with a `Sitemap:` line, and submitted in Search Console.

Remember that robots.txt is not a way to keep a page out of Google. Use `noindex` or a password for that. Our post on what to put in an XML sitemap goes into detail.

5. A canonical URL on each page

The same page can often be reached at more than one address: with and without `www`, with tracking codes, on the tool's own preview domain. A canonical tag tells search engines which one you mean. Google's duplicate URL guide treats it as a signal, not a command, recommends absolute URLs over relative ones, and says the best place for it is the HTML source.

- ✓ Every page has `<link rel="canonical" href="https://www.yoursite.com/page">` pointing to itself.
- ✓ The preview or tool domain redirects to your real domain, or is kept out of the index.
- ✓ Sitemap, internal links and canonicals all use the same form of each URL.

6. Schema and social images

Schema, structured data in JSON-LD, describes the page to machines: your organisation, a product, a set of FAQs, a breadcrumb trail. Google's structured data policies say it must be a true representation of the page, must not describe content readers cannot see, and does not guarantee a rich result. Read any schema a tool

generated. Check every rating, address and founder name against the facts, and delete what is not true.

Social images decide what shows when someone pastes your link into Slack, LinkedIn or a chat app. The Open Graph protocol lists four required properties: `og:title`, `og:type`, `og:image` and `og:url`. Give each page its own image where you can. Our free schema generator writes JSON-LD, and our post on social images per page shows a simple system.

A fix order you can follow

#	FIX	PROOF IT WORKED
1	Unique title and description on every page	View source on 3 pages shows 3 different titles
2	One H1, headings in order	Each page has one `<h1>` in the source
3	Real links in the menu and footer	Source shows `<a href` for each item
4	Sitemap on the real domain, named in robots.txt	Search Console reads the sitemap without errors
5	Canonical on every page, preview domain handled	URL Inspection shows your URL as the Google-selected canonical
6	True schema where it fits	Rich Results Test reads it with no errors
7	Social image per page	Link pasted into a chat app shows the right image

DECISION RULE

If your tool can only set titles and descriptions with a script after load, that is a rendering problem, not an SEO settings problem. Go to chapter 5.

TAKEAWAYS

- ✓ Titles and descriptions per page come first. They change what people see in search today.
- ✓ Every basic must be in the first HTML, not added by a script after load.
- ✓ One H1 per page. Pick heading tags for meaning and change size in CSS.
- ✓ A sitemap of real URLs, a robots.txt that lets crawlers in, and a canonical on every page.
- ✓ Generated schema must be true to the page. Check every fact and delete what is not.

CHAPTER 4 OF 6

Speed and phones

04

Core Web Vitals targets, why JavaScript-drawn pages load late, images, and phone testing.

The short answer

Test speed on a real mid-range phone and with PageSpeed Insights, not on your laptop. The targets are Google's three Core Web Vitals: Largest Contentful Paint within 2.5 seconds, Interaction to Next Paint of 200 milliseconds or less, and Cumulative Layout Shift of 0.1 or less, measured at the 75th percentile of page loads, with mobile and desktop counted apart.

On AI-built sites, the usual causes are the same three: the main content waits for JavaScript, images are far larger than they need to be, and too many outside scripts load at once. All three can be fixed. The first one is the hardest to fix in place.

Why it feels fast to you

You built the site on a fast laptop, on good wifi, with the files already cached from the last hundred previews. Your buyer opens a link from LinkedIn on a three year old phone, on mobile data, for the first time. Those are different tests.

- **Processor.** A mid-range phone runs JavaScript much more slowly than a laptop. A page that is drawn by script waits for that.
- **Network.** Mobile data has more delay per request. Many small files hurt.
- **Cache.** A first visit downloads everything. Your hundredth visit does not.
- **Screen.** A phone shows a narrow layout that you may never have looked at closely.

Google also uses the mobile version of your site for indexing and ranking. Its {a(MFI, "mobile-first indexing guide")} asks for the same main content, headings, structured data, titles and descriptions on mobile as on desktop, and says not to hide main content behind clicks or swipes, because Google will not trigger them.

The three numbers and what moves them

METRIC	GOOD	WHAT IT MEASURES	COMMON CAUSE ON AI-BUILT SITES
LCP	2.5 s or less	When the biggest thing on screen, often the hero image or heading, appears	Hero waits for JavaScript; huge hero image
INP	200 ms or less	How fast the page responds to taps and clicks	Heavy scripts busy on the main thread
CLS	0.1 or less	How much the layout jumps while loading	Images without set sizes; fonts and banners that push content

The LCP point matters most for rendering. web.dev's LCP guide says the LCP resource should be discoverable from the HTML source. An image added by JavaScript is found late. It also says never to lazy-load the LCP image, and to use `fetchpriority="high"` on the image most likely to be the LCP element. On a page drawn by script, the browser cannot even start on the hero until the script has run.

Images: the quickest win

Generated sites often use one large image file for every screen size. A phone then downloads an image sized for a wide desktop monitor. Fix it like this:

- **Several sizes.** web.dev's responsive images guide explains `srcset` and `sizes`, so the browser picks the right file. Three to five sizes is common.
- **Set width and height.** Give each image its size in the HTML so the layout does not jump when it loads.
- **Lazy-load below the fold only.** Use `loading="lazy"` for images further down. Never on the hero.
- **Modern formats.** WebP or AVIF are usually smaller than PNG and JPEG for the same look. Check the result by eye.
- **Text as text.** A headline inside an image is slow, invisible to crawlers and hard to edit.

Our post on image formats and sizes has more detail for Webflow sites.

Scripts: count them

Open your page source and count the script tags, then look at where each one loads from. AI-built sites often pick up extra pieces over time: an animation library for one effect, a chat widget, two analytics tools, a font loader.

- List every outside script and who asked for it. Remove the ones nobody can name.
- Load analytics and tag managers after consent, which also helps with privacy rules.
- Replace a large library used for one small effect with CSS where you can.
- Self-host fonts and load only the weights you use.

Stop the layout jumping

Layout shift is the one people notice and forgive least: they go to tap a button and the page moves. On AI-built sites it usually comes from things that arrive late and push content down.

- **Images and videos** without a set width and height, or an aspect ratio in CSS.
- **Banners** for cookies, offers or chat that slide in above the content. Put them over the page, not in its flow.
- **Fonts** that swap in at a different size. Pick a fallback font with similar widths.
- **Content drawn by script** after the first paint, such as a pricing table or logo strip that fills in a moment later. Reserve its space, or put it in the first HTML.

Testing: lab, field and a real phone

Use three kinds of test. Each one shows something the others miss.

TEST	WHAT IT IS	USE IT FOR
PageSpeed Insights, field data	Real-user data from the Chrome UX Report over the previous 28 days	How real visitors experience the page
PageSpeed Insights, lab data	A Lighthouse test in a simulated environment	Finding causes and checking a fix right away
A real mid-range phone	You, holding the phone, on mobile data	Layout, tap targets, forms and how it feels

Google's PageSpeed Insights docs say that when a page has too few real-user samples, it shows data for the whole origin, and when the origin has too few, it cannot show field data at all. That is normal for a new site. Use the lab test and the phone until real data builds up.

On the phone, do these by hand: open the home page from a link in a chat app, scroll every page, open and close the menu, send every form and check the lead arrives, and turn the phone sideways once.

DECISION RULE

If your LCP element is text or an image that only appears after JavaScript runs, image and script fixes will help but will not solve it. That is the same rendering problem as chapter 2, and chapter 5 decides it.

TAKEAWAYS

- Targets: LCP 2.5 s or less, INP 200 ms or less, CLS 0.1 or less, at the 75th percentile.
- Test on a real mid-range phone on mobile data, plus PageSpeed Insights. Not on your laptop.
- Give images several sizes, set width and height, and never lazy-load the hero.
- Count your outside scripts and remove the ones nobody can name.
- If the hero waits for JavaScript, the real fix is rendering, not tuning.

Fix in place, render on the server, or rebuild

05

Three paths, a decision table, Webflow vs Astro for the rebuild, and what a rebuild keeps.

The short answer

You have three paths. **Fix in place** when your pages already render as HTML and the gaps are titles, sitemap, images and scripts. **Render on the server** when the content is good but drawn by JavaScript, and your tool or framework offers server rendering you can switch to. **Rebuild** when rendering, editing and speed are all problems at once, or when every copy change needs a developer or a new prompt.

For a marketing site, the rebuild is usually in Webflow if your team wants to edit visually and publish without code, or in Astro if you want the most speed and control, with a CMS such as Sanity for editing. Either way, the design you like stays.

Path 1: fix in place

This is the cheapest path, and often enough. It fits when chapter 2 showed your words in the first HTML.

- Write a title and description for every page, in the first HTML.
- Add or repair the sitemap, robots.txt and canonical tags.
- Resize images and set their width and height.
- Remove extra scripts.
- Turn button-style navigation into real links.

The risk is drift. Each prompt can undo an earlier fix. Keep a short checklist from chapter 3 and run it after every change. If you find yourself re-fixing the same thing, that is a sign for path 3.

Path 2: render on the server

This fits when the content and design are right but the first HTML is empty. You keep the tool and change how pages are served.

- **Lovable.** Its SEO page says new apps already use server rendering, and older React and Vite projects can be upgraded to TanStack Start. Publish again after upgrading, then rerun the chapter 2 test.
- **v0 and other Next.js projects.** Check that marketing pages do not fetch their main content in the browser after load. Move that data fetching to the server.
- **Exported React apps.** Moving to a framework that renders on the server is real development work. Price it against a rebuild before you start.

Google's JavaScript SEO guide still calls server-side rendering or pre-rendering a great idea, because it makes the site faster for people and crawlers, and because not all bots can run JavaScript. Server rendering fixes the reading problem. It does not fix editing. Your team still changes copy through the tool.

Path 3: rebuild

A rebuild means the pages are built again on a platform made for marketing sites. It is not a redesign: the look stays. What changes:

- Every page is HTML with the words, links and headings in it.
- Copy, blog posts, case studies and FAQs move into a CMS your team edits.
- Titles, descriptions, social images, schema and the sitemap are set per page.
- Forms, analytics and your CRM are connected again and tested.
- Every URL that changes gets a 301 redirect (chapter 6).

It costs more up front than a fix. It costs less over a year if your team edits the site often, because nobody has to prompt or code a copy change again.

The decision table

WHAT YOU FOUND	TEAM EDITS COPY OFTEN?	BEST PATH
Words in first HTML, a few SEO gaps	No	Fix in place
Words in first HTML, a few SEO gaps	Yes, and it is painful	Rebuild
Empty first HTML, tool offers server rendering	No	Render on the server
Empty first HTML, tool offers server rendering	Yes	Rebuild, or render now and rebuild later
Empty first HTML, no server rendering option	Either	Rebuild
Slow on phones even with good HTML	Either	Fix in place first, rebuild if the main scripts cannot go
Site has logins, accounts or saved data	Either	Split: marketing pages rebuilt, app handled as app work

DECISION RULE

Count how many of these are true: content not in the first HTML, titles set by script, copy changes need a prompt, phone LCP over 2.5 seconds. With one, fix it. With three or more, rebuilding is usually cheaper than patching.

Webflow or Astro for the rebuild

Both publish pages as HTML. Pick by who edits and what the site needs to do.

QUESTION	WEBFLOW	ASTRO (WITH A CMS SUCH AS SANITY)
Who edits?	Marketers, visually, on the page	Marketers in a CMS, developers in code
Hosting	Managed by Webflow	On your own account with a host of your choice
JavaScript	Webflow's site scripts, plus what you add	Server rendering first, with less client-side JavaScript by default
Redirects	301 redirects in site settings, on a paid plan	In your host's config file
Custom code and data	Custom code embeds, CMS, integrations	Any code, any data source
Best for	Teams that want to change pages without a developer	Teams that want top speed and full control

If you are unsure, read our Webflow vs Astro comparison and Webflow vs AI site builders. Both are good homes for a site that started in an AI tool.

What a rebuild costs and takes

Our rebuild prices follow our normal prices: a single landing page is $500, and a website of up to 5 pages starts from $2,500. Larger sites start from $3,000. First-time clients get $500 off projects of $2,000 or more. You get a fixed quote before we start, and 30 days of support after launch for safe changes and fixes.

To get a useful quote from anyone, bring:

- The live URL and the list of pages (your sitemap or a crawl).
- Your chapter 2 results.
- Which parts must stay exactly as they are, and which can change.
- Forms, integrations and tracking that must keep working.
- Who will edit the site after launch.

See the AI-built site rebuild page for how we run it, or MVP development if your site is really an app.

TAKEAWAYS

- Fix in place when the words are already in the first HTML and the gaps are SEO basics.
- Render on the server when content is right but drawn by JavaScript and your tool supports it.
- Rebuild when rendering, editing and speed are all problems, or every copy change needs a prompt.
- Webflow suits teams that edit visually. Astro suits teams that want top speed and full control.
- A rebuild keeps the design. It changes how pages are built, edited and served.

CHAPTER 6 OF 6

Moving to the new site without losing traffic

06

URL list, redirect map, launch order and the first weeks in Search Console.

The short answer

Before launch, list every URL the old site has, map each one to its new address, and set a permanent 301 redirect for every URL that changes. At launch, publish the new site on the same domain, check the redirects, and submit the new sitemap in Search Console. After launch, watch Search Console for errors and keep the redirects for at least a year.

That is Google's own site move guide in short. It also says to expect temporary changes in rankings while Google recrawls, and that small and medium sites should move all their URLs at once. Our Site Move Playbook covers moves in full detail. This chapter is the short version, with the traps that are special to AI-built sites.

Step 1: list every URL

An AI-built site can have more addresses than you think. Gather them from every source, then remove duplicates.

- **The sitemap**, if there is one.
- **Search Console**: the Pages report and the Performance report, which show URLs Google knows about and that get clicks.
- **A crawl** of the live site. On an older Lovable app, a crawler that is not verified may get the single-page app, so it may find few links. Use the other sources too.
- **Analytics**: landing pages with visits in the last year.
- **Backlinks**: pages other sites link to. Our post on keeping backlinks in a migration shows how.
- **The tool's preview domain**, if it was ever public or shared.

Watch for hash routes such as `/#/pricing`. Google's JavaScript guide says it may not resolve fragments reliably, and a server never sees the part after `#`, so you cannot redirect those addresses on the server. If they were shared, a small script on the new home page can send them on.

Step 2: map old to new

Put the list in a sheet with two columns: old path and new path. Send each old URL straight to its closest new page. Google's site move guide says to avoid chains, because each extra hop adds delay.

OLD PATH	NEW PATH	NOTE
/pricing	/pricing	Same path. No redirect needed.
/blog/post/12	/blog/planning-treats-for-offsites	New readable slug
/about-us	/about	Renamed
/old-promo	/	No match. Home only if nothing closer fits.

Keep paths the same wherever you can. Each path you keep is one less risk. Our free redirect map tool turns your list into a Webflow CSV, a vercel.json file or a _redirects file. Our redirect checker tests them after launch.

DECISION RULE

Point every old URL at the page a visitor would want. Sending everything to the home page throws away what those pages earned.

Step 3: set the redirects

How you add them depends on the new platform.

- **Webflow.** Webflow's redirects help page says 301 redirects need a paid site plan or a paid workspace plan. It recommends no more than 1,000 rules as good practice, and wildcard rules to keep the count down. Redirects for localized paths need their own rules.
- **Webflow CSV import.** Its import page warns that a new import overwrites all existing redirects. Export first, add your rows, then import, and publish for the rules to take effect.
- **Astro.** Set redirects in your host's config, such as vercel.json or a _redirects file, so they are real server redirects.

Test the status code, not the setting. The answer should be 301 or 308, and it should land on the final page in one hop.

Step 4: launch day

#	STEP	PROOF IT WORKED
1	Final copy and content on the new site	Latest pages and posts present
2	Remove `noindex` and staging passwords from the new site	Source of key pages shows no `noindex`
3	Point the domain at the new host	Domain loads the new site on phone data and wifi
4	Redirects live	Ten old URLs from the sheet each return one 301 to the right page
5	robots.txt and sitemap live	Both load on the real domain
6	Search Console verification kept	Property still verified

#	STEP	PROOF IT WORKED
7	Submit the new sitemap	Search Console reads it
8	Forms and tracking	A test lead arrives; a test event shows in analytics
9	Rerun the chapter 2 test	Hero sentence in view source on three page types

Google's guide says to verify both the old and new sites in Search Console, including every variant (www, non-www, http, https), and to keep your verification method in place on the new site. If you used an HTML file or meta tag to verify, carry it over. Lovable's docs also say to verify a new custom domain in Search Console again after you change it.

Changing domain too?

Many AI-built sites run on a tool subdomain first, then move to a real domain. If the move also changes the domain, there is one more step. Google's site move guide says the Change of Address tool in Search Console is for moves from one domain or subdomain to another. It is not needed for http to https, www to non-www, or path changes on the same domain.

- ✓ Redirect every old domain URL to its match on the new domain, not to the new home page.
- ✓ Use the Change of Address tool once the redirects are live.
- ✓ Update links you control: social profiles, email signatures, directory listings and ads.

The first weeks after launch

Expect some movement. Google's guide says rankings may change for a while as it recrawls and reindexes. Watch for real problems, not normal noise.

- ✓ **Day 1 to 3.** Check the Pages report for new 404s. Add a redirect for each one that had value.
- ✓ **Week 1.** Inspect your top ten pages with URL Inspection. Each should show your new URL as the Google-selected canonical.
- ✓ **Week 2 to 4.** Compare clicks and impressions per page with the weeks before launch. Look for single pages that dropped, not the total.
- ✓ **Ongoing.** Search Console may warn about redirects in the old sitemap. Google's guide says you can ignore that. Keep redirects for as long as you can, generally at least a year.

Then check what AI assistants can read. Our free AEO checker looks at a page the way an answer tool would. Our post on Search Console quick wins shows what to fix next.

TAKEAWAYS

- List URLs from the sitemap, Search Console, a crawl, analytics, backlinks and the preview domain.
- Map each old URL straight to its closest new page, in one hop, with a 301 or 308.
- On Webflow, a CSV import overwrites existing redirects. Export first.
- Verify every variant of the site in Search Console and keep your verification method.
- Expect some movement. Keep redirects for at least a year and watch page by page.

QUESTIONS

Questions people ask.

Is it bad for SEO to build a site with an AI tool?

Not by itself. What matters is the output: whether the words are in the first HTML, each page has its own title, and the site is fast on phones. Some tools now render on the server by default. Run the 20 minute check in chapter 2 to see where your site stands.

Can Google read a site that is drawn by JavaScript?

Often yes, in a later rendering step. Google's own guide still recommends server-side rendering or pre-rendering, because not all bots run JavaScript. Use the URL Inspection tool in Search Console to see what Google rendered for your pages.

Can ChatGPT and other AI assistants read my Lovable site?

It depends on how the page is served. A December 2024 study by Vercel and MERJ found that none of the major AI crawlers it checked rendered JavaScript. Lovable says apps created from 13 May 2026 use server rendering, and older apps serve a pre-rendered copy to verified crawlers. Keep every sentence that matters in the first HTML.

Do I have to rebuild, or can I fix my site in place?

If your words are already in the first HTML, fix titles, the sitemap, images and scripts in place. If the first HTML is empty and your tool offers server rendering, switch to it. Rebuild when rendering, editing and speed are all problems at once. Chapter 5 has the decision table.

Will I lose my design if I rebuild?

No. A rebuild keeps the look you made and changes how the pages are built, edited and served. Spacing, type and colours are set as reusable styles so new pages match.

Webflow or Astro: which is better for a rebuilt AI site?

Webflow suits teams that want to edit pages visually and publish without a developer. Astro suits teams that want top speed and full control, with a CMS such as Sanity for editing. Both publish pages as HTML.

What does a rebuild cost at Kookies?

A landing page is $500 and a website of up to 5 pages starts from $2,500. Larger sites start from $3,000. First-time clients get $500 off projects of $2,000 or more, and every quote is fixed before we start.

My AI-built site has logins and user data. Is this guide for me?

Only for its public marketing pages. Logins, accounts and saved data are app work, where security and data matter more than search. See our MVP development page for that.

Keep going

- AI-built site rebuild
- Check if Google can read a Lovable, Bolt or v0 site
- You built your site with AI: what is left
- Can AI crawlers see your site?
- Redirect map tool
- Meta tag tool

Check what AI assistants can read with the free AEO checker

Kookies. www.kookies.io. hello@kookies.io