vinvu software

2027 EDITION

A FREE GUIDE FROM VINVU

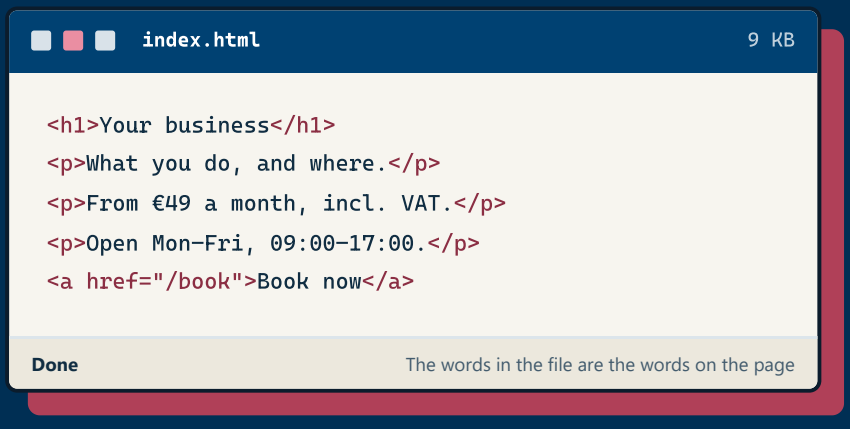
Back to the 90s

> A website that loads instantly, stays accurate, and gets found by search engines and AI assistants.

01
One source of truth

02
Fast pages, clear facts

03
Usable without our tools



Vinvu Software B.V.
De Kolk 6
8255 PE Swifterbant
The Netherlands

Contact
+31 (0) 321 700 246
contact@vinvu.nl
vinvu.nl

Registration
KvK 84836776
VAT NL863402021B01

OVERVIEW

What's in this guide

01 **Why platform sites fall behind.** 3
Where your time, your speed and your answers tend to go.

02 **The method.** 4
One facts file, pages generated from it, and a short list of rules to keep it that way.

03 **One source of truth.** 5
Why the facts file is the part worth your care, for visitors, search engines and AI assistants.

04 **Pages generated from the facts.** 7
What every page should include, and how to replace the site you have without losing visitors.

05 **What to hold the line on.** 9
Speed, indexing, AI readability, accessibility and privacy, and why each one helps you.

06 **Publish, then keep it true.** 12
What "done" looks like, and how a change travels from the file to the live site.

07 **When to choose something else.** 14
The sites this method does not suit, so you can decide with confidence.

08 **If you want this without operating it.** 15
How we run the same method for clients, and two optional tools you can add yourself.

09 **Glossary and index.** 17
Short meanings for the terms in this guide, and the pages where you will find them.

CHAPTER 01

Why platform sites fall behind

A website from the 1990s was a set of files. The words on the screen were the words in the file. A browser showed them, and a search engine read the same words. There was no page builder, no plugin stack, and no script that had to run before a price appeared.

Since then, most business sites have picked up a theme, a page builder, a slider, a chat bubble, a cookie banner and a dozen plugins. The page may look finished, but your visitor is still waiting, and so is the crawler. An AI assistant, answering from pages it has already fetched, often never reaches your price, your opening hours, or whether you have room this week.

You pay for a website so that people, and now AI assistants, can choose you. On WordPress, Webflow, Wix and similar platforms, a large share of that money goes into keeping the platform running: plugin updates, security patches, hosting, and content that has drifted between a homepage, a PDF and an old email.

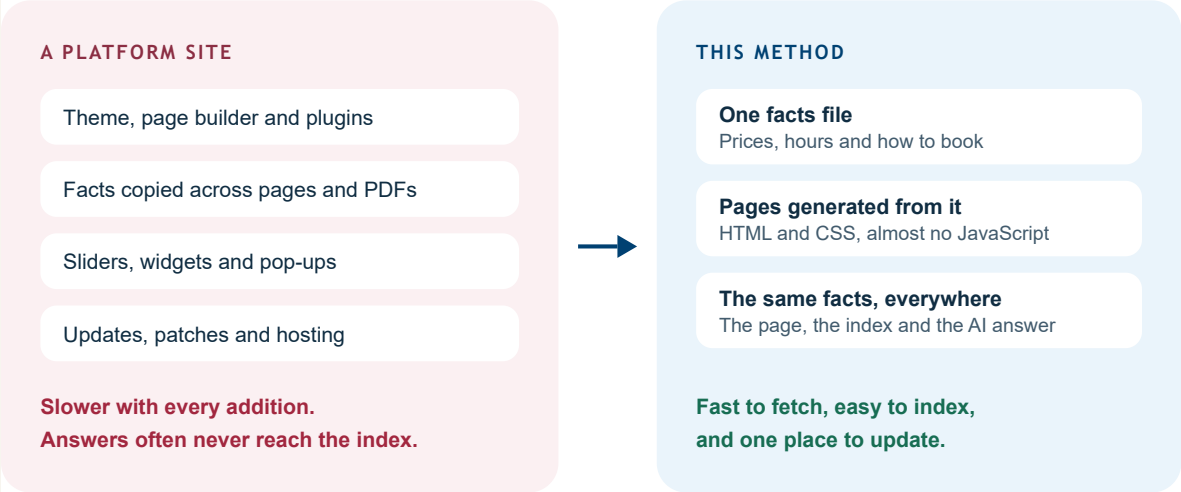


Figure 1. A platform site, and the method that replaces it.

CHAPTER 02

The method

Three things replace that stack. You can do any of them without a product from us.

One facts file. Everything a customer needs in order to choose you, written once, in a structured file. Services, prices with or without tax, discounts, terms, availability, opening hours, location and service area, how to book, the questions people actually ask, and proof you are willing to publish. Slogans can stay out. Where a source is vague or two sources disagree, the file records a question instead of a guess, so nothing uncertain reaches your customers.

Pages generated from that file. HTML and CSS, with almost no JavaScript. A static site generator is the practical way to do it; we use Astro, and plain HTML is enough for a very small site. The published result is files, so there is no CMS for you to patch. Every page reads its content from the facts file, so a price changed in the file changes on every page, and in the structured data, together.

Rules for anything added later. Your site stays fast as long as every addition earns its place. Ask an AI assistant or a developer to extend the site rather than rebuild it on a heavier framework, to never invent a fact, and to state the size of any script, font or library before adding it.

Those three pieces are really all there is to it. The rest of this guide helps you with what to put in the file, what "fast and indexable" means in practice, and how a change can move from the file to the live site.

CHAPTER 03

One source of truth

The facts file is the part that makes the site valuable after launch. On a typical site the same answer lives in too many places: a heading, a different sentence on the services page, a number inside an image, a PDF in the footer, and a note that never got published. When the price changes, someone has to go hunting.

Write the file before the design. A short excerpt looks like this:

```
services:
  - name: Indoor storage units
    summary: Dry, frost-free indoor units from 3 to 20 m², rented per month.
    options:
      - name: 5 m² unit
        price: 52.50
        available: false
        note: Expected available again in December
pricingRules:
  - title: 12-month discount
    text: 10% discount on the monthly price with a 12-month contract.
```

It works best when you keep it to plain facts. Prices are numbers, with the currency and the VAT treatment set once. If you are unsure about something, it can wait on an open-questions list. Invented reviews, invented numbers and adjectives such as "best" are better left out, because the site only says what the file says, and that is what keeps it trustworthy.

You approve the file in plain language, as a document with sections for the business, the offer, the terms, the hours, how to book, the questions and anything still open. That approval is your content contract, and design comes after it. That way, every sentence on the finished site traces back to a line you accepted.

Keeping a single source gives you four things:

- **The site agrees with itself.** The price on the homepage is the price in the table.
- **Search engines receive the same fact.** Structured data (JSON-LD for the business, the address, the hours, the services, the prices and the FAQs) is generated from the file, so it matches the visible page. Rich results become possible. They are never guaranteed.
- **A change is one edit.** "The 5 m² unit is available again" or "prices go up 5% in January" updates every page that uses it, and the structured data with them. Set a `lastUpdated` date when you do.
- **AI assistants can repeat what they can check.** They quote a fact that is visible as text. A fact that lives only in a PDF, an image, or your head is one they tend to skip. Because the page and the index draw from the same file, the next time they fetch the page they can quote your new figure.

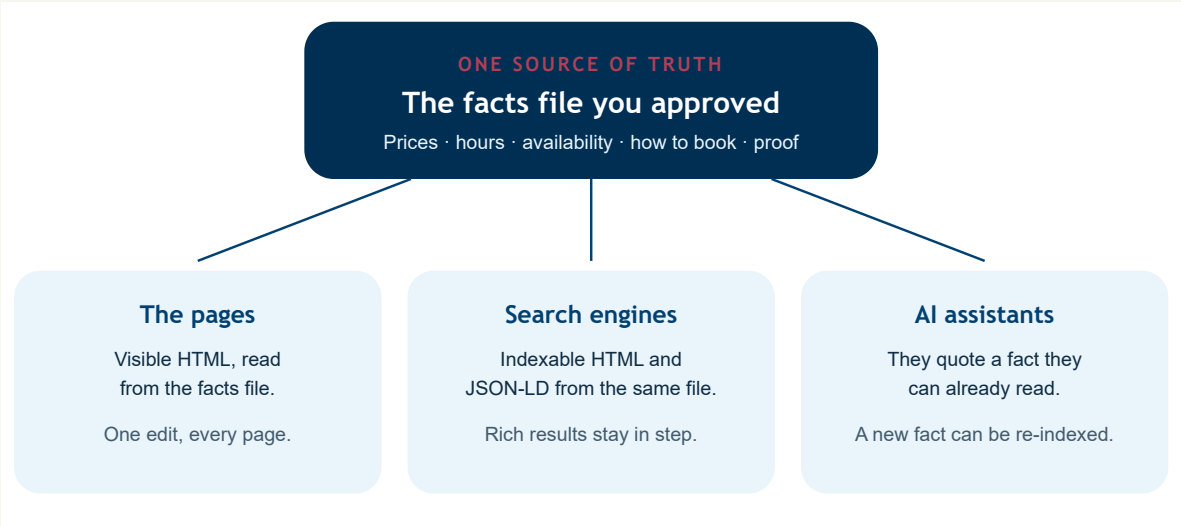


Figure 2. The facts file feeds the pages, the search index and AI answers.

Two checks are worth running before you build, and again after a big update.

The customer test. Give an AI assistant nothing except the facts file and five real questions: a price, availability or timing, a comparison, how to book, and the conditions or cancellation terms. Any answer it cannot give completely points to a missing fact. You can close those gaps with a number, a range, a from-price, or an honest "we do not publish this". "It depends" and "contact us" are the phrases that tend to cost you the mention.

The competitor gap. Look at a handful of competing sites and note whether they publish prices, availability, hours and the steps to book. Whatever they hide and you state in readable text is an advantage for you in search results and in AI answers. Put that fact where a visitor sees it without an extra click.

Prompts for this step: **1.1–1.5** and **2.1**.

CHAPTER 04

Pages generated from the facts

Once the file is approved, the build itself is ordinary. Generate the pages rather than hand-copying the facts into each template. A price typed in by hand will drift the first time the file changes, and generating saves you that trouble.

The easiest way to start is our free **Back to the 90s starter** (vinvu.nl/guides/website-starter-template). It is an Astro project with the method already in place: one facts file (`business-facts.md`), the pages generated from it, the structured data, the sitemap and `robots.txt`, and the rules for AI coding assistants in `AGENTS.md` and `CLAUDE.md`. You replace the example business with your own facts. Any other setup that follows this chapter works just as well.

Whoever builds it, this is what your site should already include:

- one clear `h1` per page, logical headings, a unique title and description, and a canonical URL
- the words in the HTML, not loaded later by a script
- JSON-LD generated from the facts file
- a sitemap that lists every page, and a `robots.txt` that points to it
- an accessible contact form: visible labels, a spam trap, and the phone number and email on the page until a form service is connected
- images with width and height, a modern format such as WebP or AVIF, and a useful alt text
- permanent redirects for any address you are replacing

It helps to give an AI coding assistant those rules in writing, in the project, so they apply on every later change. Left to their defaults, these tools tend to reach for a heavy framework and a widget. The instruction that matters most is: extend this site, state the size of anything new, and stop to ask when a requested sentence is not in the facts file.

If you are directing the build, your role is the same one we take with clients. You decide the page list and what each page should lead with. You supply colours, a logo and real photographs. You can point at sites whose feel you like and borrow just the feel. The facts always stay yours.

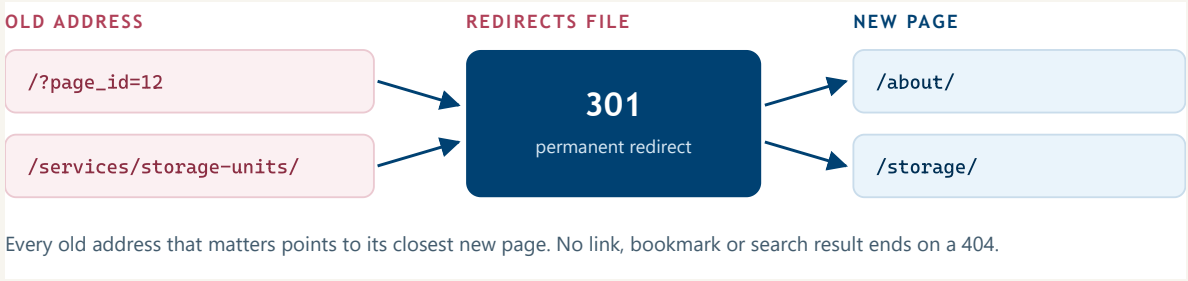
Prompts for this step: **3.1–3.5**.

Replacing the site you already have

Three steps protect the visitors and the rankings you have already earned.

- **List every old address.** Read the current sitemap. Mark the pages that receive search traffic, from Google Search Console if you have it.
- **Point each one somewhere.** Every old URL gets a permanent redirect, an HTTP 301, to the closest new page. On Cloudflare that is a `_redirects` file. On another host, use whatever that host applies at the edge. That way, an old link does not end on a 404.

- **Bring useful writing with you.** Guides and articles that earn visitors can move over as pages. Thin or outdated pages are better removed, with a redirect to the nearest useful page.



Every old address that matters points to its closest new page. No link, bookmark or search result ends on a 404.

Figure 3. Old addresses redirect permanently to the closest new page.

Switch the domain once those redirects are ready, and copy the mail records first (MX, SPF, DKIM, DMARC) so your email keeps arriving.

Prompt for this step: **3.6**.

CHAPTER 05

What to hold the line on

Four properties make your site worth more than a light template: it is fast, it can be indexed completely, an assistant can quote it, and people can use it.

Speed

An empty page scores well on its own. The real work is keeping that score once the photographs, the type and the real requests arrive. It helps to hold a budget, and to keep a lighter alternative ready for each common request.

Item	Budget	Lighter alternative
JavaScript in the browser	None by default	HTML for the menu, native `<details>` for FAQs, a normal form
Fonts	The visitor's own fonts, or one file you host	No Google Fonts or other font services
Photographs	Responsive sizes, modern formats	The main photo first, the rest lazy
Video and maps	Nothing embedded on load	A still image that links out
Analytics	None until you ask	One small cookieless script, deferred, a few kilobytes
Chat widgets and pop-ups	None	A phone number and a contact form

When you add something, note its size and what it does to the page before it goes live. A chat widget can cost more speed than the rest of the page, while a phone number and a form answer the same need.

Check your live site in PageSpeed Insights (pagespeed.web.dev) on mobile. A result in the mid-to-high 90s, up to 100, is the range this method is built to hold with real content.

Indexing

Crawlers read the HTML they are given. They are slower to fetch a heavy page, and they miss a price that appears only after JavaScript runs, inside an image, behind a click, or in a PDF. A fast page of finished HTML is indexed sooner and indexed complete, and that is what an assistant later draws on.

A few practical checks:

- View the page source and find the price, the hours and the town in the HTML itself.
- Confirm `/sitemap.xml` lists every page on the real domain, and submit it in Google Search Console.

- Confirm `robots.txt` allows the crawlers you want. That includes ordinary search engines and OAI-SearchBot, which powers ChatGPT search. GPTBot, which collects pages for training, is a separate decision. If the host has a bot-fight setting, check that it is not blocking those crawlers.
- Run Google's Rich Results Test (search.google.com/test/rich-results) and see that the structured data matches the visible text.

You can safely skip the tricks. Hidden text, invented FAQs and a special file in the root of the site are not a ranking system. `llms.txt` is a community convention some documentation sites use. It is not a formal standard, and it does not replace facts written on the page.

What an AI assistant can pick up

An assistant answering a customer uses pages it has already fetched and indexed. It can quote a business whose page states the place, the price, the discount and the availability as text. Businesses that leave those out tend to end up as "price on request".

Ask where to store Christmas decorations in Swifterbant and ChatGPT lists three storage businesses. Swift Storage comes first, with its monthly price, its discount and the note that the smaller units are full. We built that site with this method, and we have an interest in it: Vinvu has a stake in Swift Storage, and we share a building in Swifterbant.

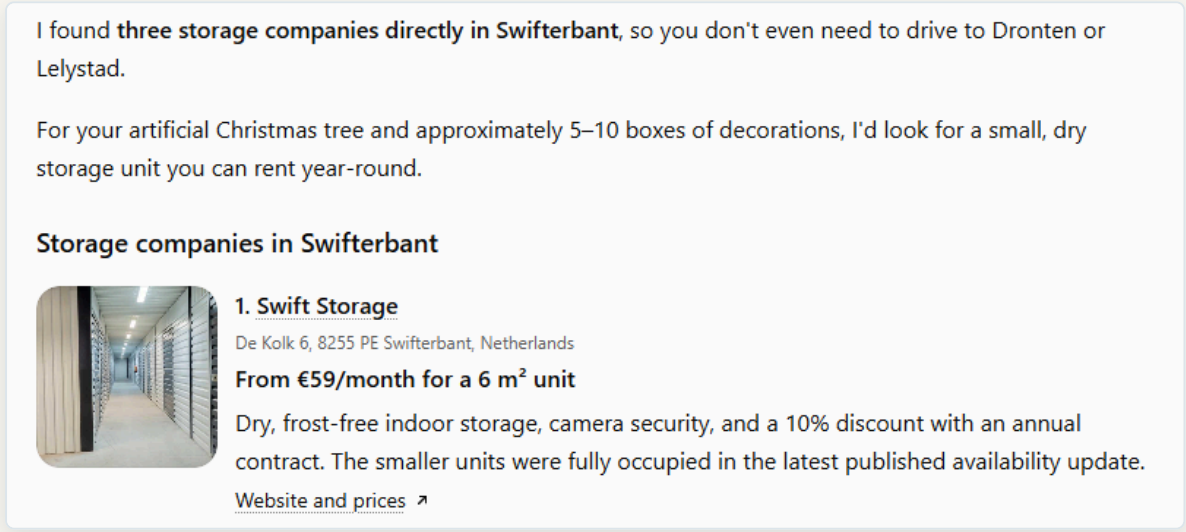


Figure 4. ChatGPT's answer to a local storage question, November 2026. Swift Storage, a business Vinvu has a stake in, is the only result with a price, a discount and availability, because its site publishes them as text.

Swift Storage also runs a booking system and a layout viewer. The public pages stay static and fast, and the booking is a small separate service beside them. Most business sites never need that split. If yours does, for live availability or a customer area, you can keep the marketing pages static and add a service only for that part.

A mention in an AI answer cannot be promised. That example is one local question in a market with limited competition. What you can build is a page a crawler can fetch quickly, with the facts an assistant needs already in the HTML, kept current from one file.

Every few months, ask the questions your customers ask, in ChatGPT, Google and one other assistant. Note whether you are named and whether the price and hours are the current ones. If a competitor is quoted instead, the gap is usually a fact your page does not yet state in text, which is good news, because you can add it.

Accessibility and privacy

Since 28 June 2025 the European Accessibility Act requires many businesses that sell to consumers in the EU online to meet accessibility requirements. WCAG 2.2 AA is the benchmark people build to: semantic HTML, labelled fields, sufficient contrast, a visible focus, a skip link, and alt text. Microenterprises that provide services, under 10 staff and turnover of no more than €2 million, can be exempt. The detail depends on your business. An accessible page is worth building either way, because more of your visitors can use it. This is general information, not legal advice.

On privacy, if you ship no tracking cookies and no third-party scripts, there is nothing on the page to ask consent for. Analytics, if you want them, are best kept cookieless and small, and named in the privacy statement. Whether a notice is still required is your own assessment, especially once an embed is added. A cookie-based tool or an embedded video can bring a consent banner back, and the banner itself is more JavaScript.

Prompts for this step: **4.1–4.3**.

CHAPTER 06

Publish, then keep it true

A static host serves the files. Cloudflare (cloudflare.com) is the straightforward choice: static requests are free on the current plans, HTTPS is included, and a push can publish the site. Workers with static assets is what Cloudflare recommends for new projects; Pages still works. Any other static host is fine if it supports your redirects and a custom domain.

You are ready to call it live when:

- the site loads on your domain over HTTPS, with HTTP redirected to HTTPS
- PageSpeed Insights shows roughly 95 to 100 on mobile
- the Rich Results Test finds structured data that matches the page
- old addresses redirect
- `/sitemap.xml` lists every page and is submitted in Search Console
- a test message arrives from the contact form, or the page shows a phone number and an email until a form service is connected

The form is plain HTML. Point its `action` at any service that accepts a normal form post. Keep the labels and a honeypot field. It is best to add no iframe and no form script: both cost speed and can hide the form from crawlers.

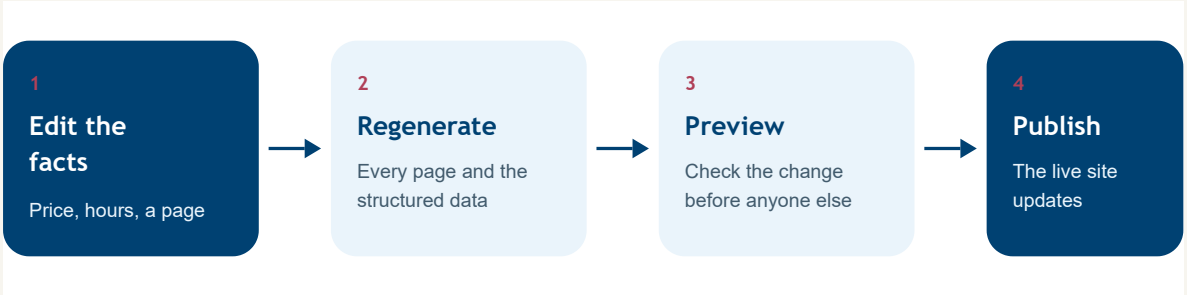


Figure 5. A change starts in the facts file, is previewed, and then published.

After launch there is no plugin for you to update. A change goes into the facts file first, including the date. Regenerate the site, check the pages and the structured data that use that fact, preview, then publish. "The 5 m² unit is available again" takes a few minutes when the file is the only place the fact lives.

Three habits keep the site earning its place for you:

- **Search Console, monthly.** Queries, indexed pages, and real-visitor speed. Add cookieless analytics beside it only if you want counts of page views and events.
- **The customer questions, every few months.** Are you mentioned, and are the quoted facts still true?

- **Prices and availability, whenever they change.** An out-of-date fact costs more trust than leaving it off. A current file is also what lets the next crawl pick up the new answer quickly.

If you do this yourself, or with a developer, someone republishes the site. That is the one remaining piece of machinery: a build, a preview, a publish. The method asks for precision more than effort, and it stays small.

Prompts for this step: **5.1–5.3**, **5.6** and **5.7**, then **6.1–6.3** once the site is live. **5.4** (a form service) and **5.5** (analytics) are optional. Formsunrise (formsunrise.com) and Fast Web Tracking (fastwebtracking.com) are two options you can choose. Any other form service works, and analytics can stay off.

CHAPTER 07

When to choose something else

This fits most small business sites: services, local companies, practices, portfolios and campaign pages.

Another tool may serve you better when:

- **You publish every day** and you want a visual editor for news or a busy blog. A separate editing layer in front of the static site can work, though a platform may be the better fit.
- **The business is a large shop.** Hundreds of products, stock and checkout belong on a commerce platform. Guides and location pages can still be this kind of site.
- **The site is an application**, with memberships, courses or customer accounts. A product built for that will serve you well, and the public pages can stay static.

Prompt for this step: **7.1**, worth running before you build.

CHAPTER 08

If you want this without operating it

The guide above is complete on its own. You do not need Vinvu, and you do not need our tools, to build or host a site this way. Any static host, any form endpoint that accepts HTML, and any AI assistant or developer who will follow the rules is enough.

What the method asks of you is care: a facts file that stays true, the patience to keep weight off, and someone who can preview a change and publish it. For many owners that someone is a developer, or they are happy to learn the publish step once. Others would simply rather not run that part, and that is fine too.

That is the part we take on when you ask us to build the site.

We design it with you. You approve the facts in plain language, then a staging preview. We publish to your domain, move the old addresses, and keep the mail records intact. The code, the build and the hosting stay with us, so there is nothing to install on your computer.

> **What it costs.** Doing it yourself costs your own time and a domain name, plus an AI coding assistant subscription if you do not have one yet. The starter is free, and static hosting on Cloudflare is free for most small business sites. Having Vinvu build the site starts at **€1,500**, and it is ready within one week.

For the sites we look after, a change can simply be a message to us. Or we connect **AI Website Manager**, our own tool for clients, between your site and the AI chatbot you already use. You describe the change in ordinary language. The assistant prepares it on a staging copy that visitors cannot see. You look at that preview, and when you say so, the live site updates. The change still starts from the same facts file, so the page, the structured data and the next index stay on one fact. There is no Git, no deployment system, and nothing to install. We set this up as part of looking after the site; it is not a product you download or connect on your own. Clients who already have it use prompt **8.1**.

Two other tools we built are optional on any site that follows this guide, including one you build yourself.

Formsunrise receives your messages. The form stays on your page as HTML, so it adds no script and stays visible to crawlers. Spam is filtered with a honeypot, submissions are encrypted and stored in the EU, and they land in one inbox. The free plan covers one website, one form and 200 submissions a month. Any other form endpoint works the same way. Prompt **5.4** covers both.

Fast Web Tracking counts page views, the steps between pages, and events such as a click on the main button. It stores counters, not people: no cookies, no visitor profiles, no session numbers. The script is about 4 KB, about 2 KB compressed, loaded with `defer`. Add it only if you want numbers, and name it in the privacy statement. Whether you also show a notice remains your assessment. Prompt **5.5** is the optional analytics step, and your site is complete without it.

If you would like us to build the site, start at vinvu.nl. Bring your current site and the questions customers ask before they call. If a platform, a shop or an application would suit you better, we will tell you.

REFERENCE

Glossary

Short meanings for the terms this guide uses. Where a term is a product, using it is optional unless the section says we set it up for clients.

301 redirect	A permanent instruction that an old address now lives at a new one, so your existing links keep working when a site is replaced.
AI Website Manager	The connection Vinvu sets up for clients whose sites we look after. A change is prepared on a staging copy, then the live site updates when the client says so. It is not a product you download yourself.
Canonical URL	The one address a page should be known by, so search engines do not treat copies as separate pages.
Facts file	The single document of services, prices, hours and the other facts a customer needs. Pages and structured data are generated from it, so you make a change once.
Fast Web Tracking	Optional analytics from Vinvu. It counts page views and events, sets no cookies and stores no visitor profiles. fastwebtracking.com
Formsunrise	An optional inbox for a form that stays as HTML on your page. formsunrise.com
GPTBot	A crawler that collects pages for model training. Allowing it is a separate choice from being found in ChatGPT search.
Honeypot	A hidden form field that people do not see. Automated spam fills it in, and those submissions can be rejected.
HTML	The text of the page as it is published. Crawlers and AI assistants read this. Words that appear only after a script runs, or only inside an image, are easy for them to miss.
JSON-LD	Structured data in the page for search engines: the business, address, hours, prices and questions. Generated from the facts file, so it matches what a visitor sees.
llms.txt	A community convention some sites place in the root folder. It is not a formal standard, and it does not replace facts written on the page.
OAI-SearchBot	The crawler used for ChatGPT search. This method leaves it allowed, so a page can be fetched for those answers.
PageSpeed Insights	Google's check of how a live page loads, including on a phone. With real content, the aim here is a score in the mid-to-high 90s.

Rich results	Search listings that can show extra facts, such as hours or prices, when structured data is valid. Possible, not guaranteed.
robots.txt	A file that tells crawlers which addresses they may fetch, and where the sitemap is.
Sitemap	A list of every page on the real domain. Submitted in Google Search Console so those pages can be found.
Source of truth	The facts file. The site shows only the prices, dates and claims that are in it, which is what keeps it accurate.
Staging	A private preview of a change that visitors cannot see. The live site updates once you accept that preview.
Static site	Pages published as finished files. There is no content system to patch between the facts file and the live page.
Structured data	Machine-readable facts that match the visible page. In this guide they are JSON-LD, generated from the facts file.
WCAG 2.2 AA	The accessibility benchmark this method builds to: clear headings, labelled fields, enough contrast, a visible keyboard focus, a skip link and a text alternative for photos.

REFERENCE

Index

Page numbers refer to this PDF.

301 redirect 7
Accessibility 11
AI assistants 10
AI Website Manager 15
Canonical URL 7
European Accessibility Act 11
Facts file 5
Fast Web Tracking 16
Formsunrise 15
GPTBot 10
Honeypot 12
HTML 7
Indexing 9
JavaScript 3
JSON-LD 5
llms.txt 10
OAI-SearchBot 10
PageSpeed Insights 9
Privacy 11
Replacing a site 7
Rich results 5
robots.txt 10
Search Console 7
Sitemap 7
Source of truth 5
Staging 15
Static site 4
Structured data 7
WCAG 11

WANT THIS WITHOUT OPERATING IT?

This guide is complete on its own. If you would rather have us build and look after the site, start at vinvu.nl.

From €1,500, ready within one week.

vinvu.nl/en/contact

The Netherlands
De Kolk 6
8255 PE Swifterbant
Flevoland

South Africa
E09 Emerald Bay
7140 Strand
Cape Town

Contact
+31 (0) 321 700 246
contact@vinvu.nl
vinvu.nl

© 2027 Vinvu Software B.V. · KvK 84836776 · VAT NL863402021B01