Engineer

Credits

Every piece of technology behind engineer.company — what reaches your browser, what builds the pages, and what proves them.

Almost everything here was made by other people and given away. This page names it, because a consultancy that will not say what it runs is asking for a trust it has not offered.

It is also the shortest honest answer to “what do you actually work with”. The list below is not a capability brochure — it is this site’s own bill of materials, with links to the projects and public checks described below.

This copy of the site is from November 19, 2026, 07:00 CET — the moment the newest page in this copy last changed. Every focused subdomain carries the same line about its own copy.

Understand a system's ingredients

The credits show the ingredients behind this site. We can make your own system equally understandable.

Explore our services Get in touch →

Ask why each dependency is there

Ask which dependencies are used, why they were chosen and how their obligations are recorded.

Cases Get in touch →

What reaches your browser

HTML, one stylesheet, local webfonts and images. **Zero JavaScript.**

HTML, written out in full before you asked for it. Every page is a file on disk, not a template assembled per request.

One CSS bundle, concatenated and fingerprinted at build time and served with a Subresource Integrity hash, so your browser can verify it arrived unmodified. Inside it rides neat-annotations by Max Syabro (MIT) — the hand-drawn arrows and margin notes, drawn in CSS alone.

The Ubuntu typeface, by Dalton Maag for Canonical, under the Ubuntu Font Licence. Served from our own machine — a webfont from someone else's CDN is a third party watching you read.

Ubuntu Mono, the same family's fixed-width cut, under the same licence and served the same way. It sets the code on these pages, and loads only on the pages that carry some.

The Shantell Sans typeface, based on the artist Shantell Martin's felt-marker handwriting, by Arrow Type / Stephen Nixon, under the SIL Open Font License. It writes the margin notes, is served the same way, and loads only on the pages that carry one.

- **Japanese lettering** uses our Engineer Japanese subsets of Noto Sans JP, © 2014–2021 Adobe, under the SIL Open Font License 1.1. The fonts are served from our own machine on Japanese pages.
- **Arabic-Indic digits** in quoted Unicode examples use small, locally served subsets of Noto Sans Arabic, © 2022 The Noto Project Authors, under the SIL Open Font License 1.1. The body, annotation and code complements load only when those digits are used.
- **The Old Italic letter KE** in Unicode examples uses locally served, one-character cuts of Noto Sans Old Italic, © 2022 The Noto Project Authors, under the SIL Open Font License 1.1.

WebP and PNG images, sized at build time. The walls these pages are built on are photographs. Two are from Unsplash: Tim Mossholder shot the dark plaster, Joe Woods the white brick. Three more are **John Snow's**, shot on a Galaxy S22 Ultra and given to this site: the pale plaster, the painted brick and the tiled wall. The light theme wears one of its four walls, drawn while the page is built, so the site is dressed afresh each time it is published. The licence does not ask for a credit; a page that names a typeface's designer should not need asking. The mark, the cog, the mascot in every pose you will meet it in, the portrait and the photographs of our home workspace in Copenhagen on the about page are ours: the desk by the window, the laptop on a table, the winter window and the balcony. They were shot or drawn here and stripped of their metadata before they were committed. The photographs in the bands across these pages are from Pixabay, under the Pixabay Content License, and each names its photographer to a screen reader. Our wallpaper collection offers artwork to download at its native resolution.

One borrowed icon set, and one drawing from a second. The pictures of what a piece of work was made of — one beside every achievement in every list, a row of them under every review, one behind each proof figure and each service heading — and every other icon this site draws: the marks on both menus and behind each page, the contact line, the networks, the language switch, the arrows and the chevrons. All of them but one come from Font Awesome Free 7.3.1 by Fonticons, under CC BY 4.0: 353 drawings in all, kept in this repository and served from this host like everything else. 51 are not drawings of an idea but the technologies' own marks — PostgreSQL, Linux, Docker, Go, Swift, Rust and the rest — naming what a job was actually built on. A mark is used where the record supports it and a plain picture where it would not: the database work also runs on MS SQL and BigQuery, so it keeps a drawing of a database rather than any one vendor's logo, and work built away from something never wears that thing's mark. A trademark stands for its owner, so an exact borrowed mark is better than one redrawn to match our set, and where there is no mark to borrow we use a plain picture instead of inventing one. The rows of logos that say which products we work with use the same marks, and each stays its owner's trademark: naming a product claims no endorsement from its owner. The one exception is the processor chip marked AI that stands beside the picture of each of our AI services, wherever that picture appears, and beside the menu entry for AI agents. It is `cpu-ai`, the only drawing taken from Lucide 1.61.0 by Lucide Icons and Contributors, under the ISC License, and it is kept and served the same way. The identity itself is ours throughout: the brand mark, the cog, the mascot, the portrait and the office photographs. The credit sits on this page rather than in the files, because the build minifies its HTML and strips the comment each icon ships its attribution in.

Zero executable JavaScript is published. Menus and links use the browser's native HTML controls; the structured data in the page is inert JSON-LD. No cookies are set and no resources are requested from another company's server. There is no consent banner. Open the network panel and count the hosts: one. The public page totals come from the masked server log and add no browser code, tracking request or browser storage.

How the pages are made

Hugo — the static site generator, written in Go. Templates, taxonomies, every language and the feeds are its work.

Markdown and Go templates for the hand-written pages and the layouts.

SQLite — the portfolio is not authored here. It lives in a multilingual CV database and is exported into this site by a Python script, so a fact is written once and appears wherever it belongs.

Python for that exporter, and for generating the icon set with **ImageMagick** from a single drawing of the mark.

Chromium renders each page's print layout as a PDF in the local build; qpdf compresses the PDF structure before publication.

How it is proved

There is no hosted build service; the quality gate runs on the machine the site is written on, and every diagnostic is an error rather than a warning.

- **vnu** validates the HTML, **Stylelint** the CSS.

- **axe-core** audits representative pages across templates and every configured language against WCAG 2.2 A and AA rules, in both colour themes, as part of the release gate. The findings cover the automated rules exercised; the accessibility statement records the manual-review and assistive-technology limits.

- **Puppeteer** drives headless Chromium to measure contrast, focus order, target size, reading measure, viewport overflow and print pagination, and to compare visual output through **pixelmatch**. Contrast, focus, target-size, reading-measure and print checks run in the release gate; viewport overflow and visual regression have on-demand runs. Puppeteer also generates the share card from the site's own stylesheet and typeface.

- **Lighthouse** holds the performance budgets; **lychee** checks that every link still resolves.

Prettier, **markdownlint**, **ESLint**, **Ruff**, **basedpyright**, **yamllint**, **ShellCheck** and **codespell** keep the source honest; **gitleaks** scans the source for recognized secret patterns.

Task runs all of it under one name per job, and **pre-commit** runs the file-level half again on every commit, so the gate is not something anyone has to remember.

Where it runs

One small machine does all of it: a single-core **DigitalOcean** server with 512 MB of memory and a 10 GB disk, standing in **Toronto, Canada**, answering at **165.227.39.206**. That address is what DNS already hands to anyone who asks, so printing it here saves you the lookup rather than telling you anything new. What the machine records while it answers is on the privacy page.

Debian is the operating system.

Caddy answers every web request and renews the HTTPS certificate itself, from **Let's Encrypt**.

Agate, **Gophernicus** and **Tor** answer the other three: this site is also a Gemini capsule at `gemini://engineer.company`, a Gopher hole at `gopher://engineer.company` and an onion service at `http://engineer2ezpeji52ipphljcflzyhetvy5r6tnqubtq45nzlxxv6rsqd.onion/`. All four read the same folder of files, so one deploy updates all four and none of them can drift from the others. The onion is the only front door we could add without opening a port: tor dials out to introduction points, so the machine's open ports stay exactly as they were. The web site sends an `Onion-Location` header too, so Tor Browser offers it without being asked. It begins `engineer` because we spent the CPU hours to make it readable — a convenience, not a proof. The same prefix is available to anyone willing to spend the same hours, so compare the whole address rather than the start of it.

Ansible describes that server as code, so it can be rebuilt rather than remembered.

Soft Serve is the git server behind the Git Hub at git.engineer.company. This site's own source lives there.

restic is the tool in our off-host backup and restore-test implementation. Activating its storage and verifying a live restore are separate steps.

Cloudflare answers DNS — and only DNS. The site is deliberately not proxied: openness is worth more to us than the protection, and a proxy would put one more company between you and us.

Mail is **Apple iCloud**, which is why the address on the contact page reaches a mailbox and not a marketing platform.

Where else it answers

Everything above describes one site, and it lives at engineer.company. The same work is also served at a focused address for each trade: whoever arrives at one of them is shown that trade first and the company's own pages unchanged, and every one of those pages names the copy at engineer.company as the original.

data.engineer.company — databases and the pipelines that fill them.

platform.engineer.company — cloud infrastructure and the machinery that deploys it.

software.engineer.company — software built to be maintained rather than only delivered.

advisory.engineer.company — consulting, architecture and technical leadership.

brand.engineer.company — brand and marketing work.

support.engineer.company — the machines, networks and helpdesk an office runs on.

ditte.engineer.company — the founder's own pages, which also answer at daryna.engineer.company.

And **git.engineer.company** is the front door of the git server named above: the repositories are private and reached over SSH, so that page is all the web is shown.

The site is also complete in **Ukrainian**, **Norwegian**, **Swedish**, **Spanish**, **German**, **Japanese** and **French** — with indexable translated pages listed in the sitemaps beside their Danish and English counterparts. The strip in the top bar does not offer these languages, so these links are the way in.

Standards it follows

- **WCAG 2.2 level AA** is our accessibility target across pages, languages and colour themes. The current CSS icon labels have no keyboard dismissal control, so 1.4.13 is met only in part. The accessibility statement records our tested coverage, the known limitation and how to report a barrier.

- **schema.org** structured data, marking only what is true and visible on the page.

- **Web App Manifest**, so the site can be installed as an app.

- **RSS**, **Atom** and **JSON Feed**, all three carrying the full text. Two XML feeds rather than one because some readers take only one of the two and neither is reliably the one.

- **RFC 9309** `robots.txt`, which allows everything, to search crawlers and AI systems alike — the writing here exists to be read, and it names the sitemap both of them follow. There is a readable index of the same pages, for anyone who would rather browse it than parse it.

- **RFC 9116** `security.txt`, plus `humans.txt`, `llms.txt`, `agents.txt` and `ai.txt` for the machines that come asking — and `ads.txt`, which authorises nobody, because there are no advertisements here to sell.

- **Gemini**, **RFC 1436 Gopher** and **Tor onion services**: the same pages, served over all three at `gemini://engineer.company` , `gopher://engineer.company` and `http://engineer2ezpeji52ipphljcflzyhetvy5r6tnqubtq45nzlxxv6rsqd.onion/` , out of the same folder of files this page comes from — one deploy updates all four. Gemini and Gopher each want their own client rather than a browser; the onion wants a Tor client, and Tor Browser is the ordinary one. Gopher has no encryption layer at all, so it carries which page you read in the clear; Gemini encrypts that, but the network you read from still sees which address you asked for. Over the onion it sees neither: of the four, that is true only there. The address itself is the service's public key, so the circuit is authenticated and encrypted before a word of HTTP is spoken — the `http://` is correct rather than careless, and a certificate would only prove a name the address already proves by arithmetic.

Check it yourself

Everything above is a claim, and the earlier section on how it is proved names the tools we run against our own machine — which is us marking our own homework. These are the same questions asked by people who have never heard of us. Every link runs a live test against this page, or the address bar of the one you are on; none of them takes our word for anything, and none of them is ours.

- **The W3C Markup Validator** parses this page's HTML against the living standard. It is the same engine as the `vnu` above, run by the W3C rather than by us.
- **The W3C CSS Validator** reads the stylesheet this page is wearing. Expect it to report errors, and read them before believing them: as of August 2026 it validates against a profile that stops short of `system-ui` , `@property` , `clamp()` in a `font-size` and scroll-driven animations — all four shipped in browsers, and all four used here. Against /terms/ it reported 176 complaints, and 175 were of exactly that kind. The one that was not had caught a real defect: a resting `rotate: 0deg` loses its unit to the minifier, `rotate: 0` is not an angle, and the browser drops the whole line — so an animation shipped without the rotation it was written with. Stylelint cannot see that, because it reads the stylesheet we write and the minifier rewrites it on the way to you. Stylelint remains the authority we hold the source to, and a check of our own now reads what was actually served.

The W3C Feed Validator checks the Atom feed, and the same service checks the RSS one. It warns that most entries share a modification date, and that warning is correct: the dates come from git history, and most of these pages genuinely were written in the same few commits. We would rather show a warning than invent a date.

The JSON Feed Validator checks the third feed, which the W3C service does not cover — it predates the format.

PWABuilder reads the manifest and the rest of what makes the site installable, from Microsoft's machines rather than ours.

The W3C Internationalization Checker reads the language and encoding declarations that make every language version of this site legible to the right reader.

The Schema Markup Validator reads the structured data, and **Google's Rich Results Test** checks supported markup for Google rich results.

PageSpeed Insights runs Lighthouse from Google's machines rather than ours — performance, accessibility, best practices and SEO, with real-world data where enough people have visited.

The Mozilla HTTP Observatory grades the security headers, and **Qualys SSL Labs** grades the HTTPS configuration behind them.

Security Headers is a second opinion on the first of those, from someone else entirely.

What none of them can establish on its own is full accessibility conformance, and we would rather say so than let a green badge imply otherwise. **Automated tools cannot check every accessibility requirement** — human judgement is needed to assess whether a heading describes its section or an alt text says the useful thing. The audits above are the floor, not the proof. If you use a screen reader here and something reads wrongly, tell us — that is the report we cannot generate.

Thank you

Every project above is maintained by people who did not have to share any of it. If you build on this stack too, that is where your money and your bug reports are worth most.

And if this stack reads like your kind of engineering, put ours to work.

For another view of the work behind the website, browse the site in numbers.

Separate your work from its components

Begin with a dependency inventory that distinguishes your own work from external components.

See what it costs Get in touch →

Make the handover understandable

Discuss documentation and delivery that let the next maintainer understand what they inherit.

Explore our services

welcome@engineer.company +45 50 10 61 84 Copenhagen, Denmark

Take this page with you

Contact card

Carefully Crafted With Love by D.

CVR 45969703 · Tagensvej 105C, 3. th, 2200 København N

Copyright © 2025 – present · Engineer ApS · All rights reserved.

Printed from https://ditte.engineer.company/credits/