Engineer Advisory

Privacy

How this site handles data: no cookies, no tracking scripts, and a server log that never holds a full IP address.

This site sets **no cookies**, loads **nothing from a third party**, and runs **zero JavaScript**. Its pages, menus and links use HTML, CSS and the browser's own controls. The structured data for search engines is data, not executable code.

There is no analytics tag, tracking pixel or embedded third-party script or font. Everything needed to read the page arrives as HTML, styles, local fonts and images. What is and is not stored in your browser has a page of its own: cookies, with the short answer and the way to check it.

That is not a policy we promise to keep; it is how the site is made. You can check it: view the source, or open your browser's network panel and watch which servers it talks to. The answer is one — ours.

Build with data restraint

The privacy policy explains this site's data handling. We can also help design restraint into your own systems.

Explore our services Get in touch →

Trace what your service records

Inspect what your service collects, why it is needed and where it travels.

Cases Get in touch →

What the server records

Our web server keeps an access log, the way essentially every web server does. It is the only way we can tell which pages people read, and whether anyone finds their way to the contact page. One piece of configuration writes every log on that server, so each address that records anything records it the same way, and that covers more addresses than this one: fifteen today — this site, the `www.`

version of it, our `git.` subdomain, the seven focused addresses the credits page lists (one of which answers to a second name as well), the short address that redirects to this page, `engineer.center` and its `www.` version, which redirect to this site, and the host that serves nothing but our mail policy. For each request it stores:

- **The page requested**, as your browser asked for it, and the date and time. A link can carry more after a `?`, and the log keeps that as it was sent; the summaries below keep the page alone, plus one word if the link said where it was published (`utm_source` — ChatGPT adds `utm_source=chatgpt.com` to the links in its answers).
- **Whether it worked** — the response status, so we can find broken links.
- **Your browser's self-description** (the "user agent"), which is how we separate people from search-engine crawlers. A crawler that gives a name is checked against the networks its owner publishes, using the blurred address below.
- **The site you came from**, if you followed a link — see below.
- **A blurred network address** — see below.

Two things worth being precise about

Your IP address is never stored in full. It is shortened at the moment the line is written, not cleaned up afterwards: the last part is replaced with zeroes before anything touches the disk. What remains identifies a rough network, not a person or a household — an entire office and an entire street can look identical. The server records the address under two names and masks both, and it drops the headers that can carry an address of their own — the ones a company proxy or a content delivery network puts in front of you — before the line is written. What reaches the disk is that rough network, which is not much to lose, hand over, or be asked for.

We do keep the site you came from, and that is a real disclosure. If you arrive by clicking a link, your browser tells us which page it was, and we keep the site's name — `example.com`, not the full address of the page, and never anything after the `?`. We keep it because it is the only way to know whether anyone finds us at all, and from where. We would rather write that plainly here than collect it quietly.

How long, and who sees it

A line lives about **thirty days**, on every address, busy or quiet. Our server software closes each log file once a day, and sooner if it reaches about **10 MB**, and deletes a closed file twenty-eight days after closing it. Add the day a line can wait for its file to close and the day a closed file can wait to be deleted, and that comes to thirty. Every night a separate check reads every file back and fails if any line in it is older than that.

This has gone wrong twice, and we would rather tell you about both than fix them quietly. In September 2026 we found four quiet addresses that filled 10 MB so slowly that their log had never been closed at all and still held every line since we switched logging on. The daily job we added then to close those logs did not hold. On 26 October 2026, seven of our thirteen addresses were still writing into a file the job had already marked as closed and would not delete while it kept growing, and one address still kept a closed file forty-seven days old. The server software now closes and deletes the files itself, and the job is gone.

Summaries outlive the log, and we would rather say so than let you find out. A program on the server reads the log and prints counts; the counts travel to our own computer, where the summaries are written and kept in a private repository of ours. Each address has four of them — the last thirty days for people, the same for everything automated, a row per day and a ranking per month — plus a page for each month and each year. They are written afresh on every run and the repository keeps its history, so every earlier edition of them stays too.

A daily tally, kept indefinitely: so many pages read, so many of them by crawlers, so many reaching the contact page. Dates and counts, nothing else.

A monthly ranking, also kept indefinitely: which pages were read most that month, which sites people arrived from and which source a tagged link named, which crawlers and automated scanners called, and which blurred networks the automated probes came from — each as a name and a number of times. It is how we can still answer “what did people read last March” once March’s log is gone.

What that ranking is not: it is never a record of what any one visitor did. A page has a count beside it and nothing else — not who asked for it, not when beyond the month, not what else they read. The lines are never joined back together,

because the pieces needed to join them are not kept.

We do not send the log or private reports to an analytics company. The private summaries stay on our own computer. The public page totals below are the limited part shown to any reader.

Public page counts. A page's counter estimates successful HTML GET requests from the existing masked web log, separately for each site and page. Known bots, probes, assets and failed requests are excluded. Reloading can count again; these are estimated page views, not unique people or evidence of a conversation or an order.

Recorded views come from available log history. The collector can import older retained masked logs, preserving its original initialization record and the observed log period. Older history can have gaps; known observations still count, with that limitation disclosed. This is not a complete lifetime total, and old daily or monthly report estimates are not added. The server updates the number roughly every thirteen minutes, without browser JavaScript, cookies, extra tracking requests or browser storage.

The counter keeps aggregate totals, observed time bounds and internal log-reading positions and file bookkeeping, not visitor identifiers or individual request histories. Totals can outlive the raw log. A dash means the counter is not initialized, is stale, or has lost reliable coverage of newly collected logs. Known gaps in older imported history do not erase recorded observations. A valid zero means no qualifying request was observed in the available history. Existing raw-log retention and the daily and monthly report archives remain unchanged.

The other ways in. This site is also served over Gemini and Gopher, two small protocols with no access log at all, and over Tor. Since September 2026 we count **how many requests** the Gemini and Gopher copies answered on each day, and nothing else: no address, no page, no visitor. **We count nothing at all over Tor.** Visitor counters and analytics stay disabled for our onion service. Tor is designed to hide who accesses it.

Counting them showed us something we had not known, so here it is. The Gopher copy is started afresh for each visitor by the operating system, which names that moment after **the address it came from** and writes the name into its

own log. Our server software never sees it and no setting of ours could have stopped it. We found it the day we first counted, and that log now expires after thirty days like everything else here. It was keeping things for longer.

Where our own server is. One machine, in Toronto, Canada. The log is written there and stays there; the summaries are written from it on our own computer. The credits page names the server, along with everything else this site is built on and served by.

We do not build profiles, and we cannot follow you across other websites. For as long as a line is in the log it carries no name, no email address, no cookie and no complete address. It does carry your browser's full self-description, a blurred network, the language your browser asks for and a timestamp to a fraction of a second, all on one line. A rare browser on a small network narrows that down further than we would like, so we write it here rather than tell you that identifying you is impossible.



Getting in touch

Mail sent to welcome@engineer.company is read by us and kept in our own mailbox, hosted with Apple iCloud. We use what you write to answer you. We do not add you to a mailing list, and there is no mailing list to add you to.

If you want to know what we hold about you, or want it removed, write to the same address and we will answer honestly — which, given the above, will usually be "nothing that identifies you".

Questions this page does not answer go to the same place — ask us directly.

Map collection, access and retention

Start with a map of data flows, access and retention. That makes the technical changes assessable.

See what it costs Get in touch →

Discuss a data-handling requirement

Discuss a concrete data-handling requirement for your software; legal interpretation stays with your advisers.

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://advisory.engineer.company/privacy/