

Security & Compliance Package

Vaethra · everything a vendor security review asks for

Issued 11 October 2026

The current version of this document is always at vaethra.com

security@vaethra.com · vaethra.com/trust

Contents

1. Trust centre — vaethra.com/trust
2. Security — vaethra.com/security
3. Architecture — vaethra.com/security/architecture
4. Access control — vaethra.com/security/access-control
5. Data protection — vaethra.com/security/data-protection
6. Privacy — vaethra.com/security/privacy
7. Subprocessors — vaethra.com/security/subprocessors
8. Availability — vaethra.com/security/availability
9. Business continuity — vaethra.com/security/business-continuity
10. Incident response — vaethra.com/security/incident-response
11. Vulnerability management — vaethra.com/security/vulnerability-management
12. Secure development — vaethra.com/security/development
13. Compliance — vaethra.com/compliance
14. Support — vaethra.com/support

Trust centre

Vaethra's security, availability, privacy and compliance documentation, in one place.

TCVaethra reference

The answers a vendor security review asks for, published in advance

A vendor security review usually means a questionnaire, several weeks of email and a call. The questions are largely the same each time, so they are answered here in the detail a reviewer needs: architecture, access control, availability, recovery, incident handling, vulnerability management, how code reaches production, what happens to data over its whole life, and which frameworks we work to. It is written before the first enterprise agreement on purpose: so that the first enterprise conversation is not stalled for a quarter by a questionnaire, and so that a reviewer can see what is running and what is committed before anyone is asked to trust a summary.

Two points about scope. Where a control is planned rather than running, or committed rather than yet exercised, the relevant page says so on the page itself rather than only in an appendix. And evidence — the risk register, the records of processing, access reviews, vendor assessments, the policy set, and the recovery test report once the first exercise has run — is provided under NDA rather than published, because those documents describe operational detail in a form more useful to an attacker than to a reviewer.

eighteen public pages

Security programme

PAGE	COVERS
<u>Security overview</u>	Approach, infrastructure, encryption, authentication, monitoring, logging, personnel, physical security, governance, contacts
<u>Architecture</u>	Components, trust boundaries, environments, data flows, model processing, tenancy, network boundaries
<u>Access control</u>	MFA, least privilege, production access, access reviews, joiners and leavers, credential management, service accounts
<u>Data protection</u>	Classification, the licence filter, encryption, retention, deletion, residency, ownership and export
<u>Privacy</u>	Categories of personal data, lawful bases, retention, individual rights, privacy governance
<u>Subprocessors</u>	The four subprocessors, what each holds, where, their attestations, and 30 days' notice of change

Operations

PAGE	COVERS
<u>Availability</u>	Uptime measurement, availability architecture, monitoring, maintenance, and the enterprise service level
<u>Business continuity</u>	Backups, recovery objectives, five disaster scenarios, recovery procedure and testing
<u>Incident response</u>	Classification, detection, process, the 72-hour notification, and the post-incident report
<u>Vulnerability management</u>	Dependency monitoring, patch windows, security testing, and responsible disclosure
<u>Secure development</u>	The path to production, change management, environment separation, secure coding practice
<u>System status</u>	What is running now, answered by the running system
<u>Support</u>	Channels, hours, response targets, escalation, notice periods

Compliance and legal

PAGE	COVERS
<u>Compliance</u>	Framework status, and a control-by-control map of SOC 2, ISO 27001 Annex A and the GDPR against where each is implemented
<u>Enterprise Service Addendum</u>	The availability commitment, the service-credit schedule, support targets, notice periods, residency and escrow, as an enterprise order form incorporates them
<u>Data Processing Agreement</u>	Ready to execute, with the Standard Contractual Clauses and the annex of technical and organisational measures
<u>Privacy Policy</u>	The formal notice
<u>Terms</u>	Terms of use, and how an enterprise agreement differs

Our security contact details are also published at </.well-known/security.txt> in the format defined by RFC 9116.

Evidence, available within two business days

The pages above state what our controls are. The documents below evidence that they ran, and are provided to customers and to organisations in an active evaluation under a mutual NDA. Each row says whether the document exists today or depends on an exercise that has not yet been run, because a document offered by name and not yet written is exactly the failure this page is written against.

DOCUMENT	CONTENTS
Security policies	Exists. Access control, change management, incident response, cryptography, acceptable use, vendor management, business continuity. Written 12 September 2026.
Risk assessment	Exists. The register: risk, likelihood, impact, owner, treatment. First review 12 September 2026; reviewed quarterly.
Records of processing	Exists. The Article 30 record, by category of data and purpose.
Access control evidence	Exists. The access review of 12 September 2026 — every platform, every holder, the second-factor state of each — and each quarterly review after it.
Vendor assessments	Exists. Our review of each of the five subprocessors, and their attestations.
Vulnerability assessment	Exists, live. The dependency audit (zero open findings on 11 October 2026) and the open integrity findings published on /status , with remediation dates.
Completed security questionnaire	Committed. Completed in writing within five business days of a request — CAIQ, SIG Lite or your own format. None has been requested yet.
Disaster recovery test report	After the first exercise. The written test plan exists; the first restore to a scratch instance runs before the first enterprise agreement, and the report — what was restored, from which backup, how long it took, what failed — joins the pack then.
Backup restoration evidence	After the first exercise, as above. Evidence that backups restore, not that they are taken.
Penetration test report	After the test. An independent test is to be commissioned before the first enterprise agreement; no firm is engaged yet.
Incident response exercise	After the first exercise. No incident has been declared and no tabletop has been run; the first is scheduled with the disaster-recovery exercise.
Compliance evidence	Being assembled. The control-by-control map on /compliance , with the artefact behind each row attached as it is produced.

Request it from security@vaethra.com with the entity name and the NDA you would like signed, or ours.

Three terms, used consistently

TERM	MEANING
Live	Running in production. Where you can verify it independently, the page says how.
Committed	A contractual undertaking — a response time, a notification deadline, a recovery objective. It binds us under the agreement.
Planned	Not yet in place, with the dependency stated.

Figures carry the date they were measured. Where a page states a number — tables with row-level security, functions executable by the public role, automated tests — it was taken from the running system on that date and can be re-checked on request.

If you find a statement here that does not match what you can observe, please tell us at security@vaethra.com.

The same documents as files, for the people who file things

A vendor record wants one attachment, not eighteen URLs. Counsel wants a contract that can be printed and signed. And a document read by a language model — now a normal step in a procurement review — is easier to supply as one file than as a crawl. All four carry the date they were issued and are rebuilt from the pages above, so a PDF cannot quietly diverge from what the site says.

Nothing here is under NDA. The evidence pack in §3 — policies, the risk register, recovery test reports, access-control evidence — is a separate request to security@vaethra.com.

Every address, and what it is for

SUBJECT	ADDRESS
Vulnerability reports and security incidents	security@vaethra.com
Security questionnaires, evidence requests, DPAs	security@vaethra.com
Privacy and data subject requests	privacy@vaethra.com
Contracts, terms and licensing	legal@vaethra.com
Faults, support and enterprise escalation	support@vaethra.com , and +1 253 402-5005 for an enterprise Sev-1
API access, pricing, general enquiries	contact@vaethra.com
Data providers, about our crawling	ops@vaethra.com

Security contact details are also published at </.well-known/security.txt> in the format defined by RFC 9116, so a scanner finds them without reading this page.

Legal entity and notice address

Registered name	Vaethra Technologies LLC
Form and jurisdiction	Limited liability company, State of Wyoming, United States
Mailing address	30 N Gould St Ste N, Sheridan, Wyoming 82801, United States
Notices	The address above, or legal@vaethra.com

This is the entity that contracts, holds the infrastructure accounts and signs the [Data Processing Agreement](#). It is established in the United States, so for a customer in the EEA or the UK it is the data importer under the Standard Contractual Clauses rather than the exporter.

Security

How Vaethra is built and operated: infrastructure, encryption, authentication, monitoring, personnel and governance.

TC 01Vaethra reference

Three principles the system is built around

Controls are verified, not assumed. A scheduled job can complete successfully and accomplish nothing — an upstream can return HTTP 200 with an error in the body, a job can catch its own exception and continue. Vaethra runs a set of automated checks that test outcomes rather than exit codes, including a check on the licence filter itself. Findings are recorded on every run with the transition that caused them — appeared, cleared, still open — retained for review, and the open set is published on [the status page](#).

We hold as little of your data as possible. The Terminal and Atropos require no sign-in, so there is no account, no profile and no record of what any visitor viewed. What we hold is public-source records, the sign-in addresses of the people an API customer admits to its workspace, key records, and the enquiries sent through the one form on this site. This limits the impact of any compromise more effectively than any control we could add on top.

Our position is documented, including the gaps. The [compliance page](#) lists every framework we do and do not hold, and each page below states what is running and what is planned. If you find a statement here that does not match what you can observe, please tell us at security@vaethra.com.

Take it with you

These pages are also published as one PDF, so a reviewer can attach it to a ticket, read it offline, or put it through a model without clicking through eighteen pages.

Managed infrastructure, two providers, nothing self-hosted

Vaethra operates no servers. There is no operating system to patch, no container image to maintain, and no host that accepts a remote shell.

LAYER	PROVIDER	HOLDS
API and applications	Cloudflare Workers, globally distributed	No persistent state; every instance is stateless and ephemeral
Static sites	Cloudflare Pages	Pre-built HTML, no server-side execution
Analytical database	Supabase managed PostgreSQL 17, AWS us-east-2	The event and entity record, API key digests, alert rules
Connection pooling	Cloudflare Hyperdrive	Nothing durable
Cache and coordination	Cloudflare KV and Durable Objects	Cached responses, cursors, rate-limit counters — all reconstructible
Object storage	Cloudflare R2	Pre-built bulk extracts and cached satellite scenes
Language models	Cloudflare Workers AI	Inference runs inside Cloudflare's network

The practical effect is a smaller attack surface than a self-hosted equivalent: no VPN, no bastion host, no fleet to keep patched. What remains is our application code, our configuration, and the two providers' own controls. Cloudflare maintains SOC 2 Type II, ISO 27001, ISO 27701 and PCI DSS; Supabase maintains SOC 2 Type II. Both cover the infrastructure they operate.

All data encrypted in transit and at rest

PATH OR STORE	PROTECTION
Browser to edge	TLS 1.2 minimum, TLS 1.3 preferred. HSTS for one year including subdomains. HTTP is redirected and never served.
Edge to database	TLS through Cloudflare Hyperdrive. The database refuses any connection without TLS and any connection from outside Cloudflare's published address ranges (both settings applied and measured 11 October 2026), so no other client can reach the listener with the password alone.
Database at rest	AES-256, including backups and replicas
Object storage at rest	AES-256
Secrets at rest	Encrypted in the platform secret store, injected at runtime, not readable back once set
API keys	Stored as SHA-256 digests. A copy of the database yields no usable credential.
Endpoint devices	Full-disk encryption and screen lock

The zone is signed with DNSSEC at Cloudflare; the DS record that completes the chain is entered at the registrar when the domain lock lifts, and until then a resolver does not validate it. TLS certificates are issued and rotated automatically. No certificate in the estate is managed by hand, so none can expire unnoticed.

Four classes of caller, each with its own gate

CALLER	AUTHENTICATION
API customer	A bearer key, compared as a SHA-256 digest. Displayed once at creation and not recoverable afterwards. Revocation clears the edge cache immediately rather than at the next expiry.
Terminal and Atropos	First-party origin. Both are free and account-less, and receive no access an issued key would not also give.
Administrative endpoints	Two independent gates: an identity-aware proxy at the network edge, and a bearer token checked again in the application. Neither is origin-based.
The application, to the database	Two pooled connections through Cloudflare Hyperdrive, whose credentials live in the platform's configuration and never in the repository. The public record is served on a connection that holds the database owner role. The customer plane is served on a separate connection as a role that cannot bypass row-level security, whose only route to a customer row is a transaction-scoped <code>SET ROLE</code> into the tenant or operator role.

fifteen paths (15) are open without a credential — health, status, uptime, facts, the accuracy record and the backtest ledger, the graph statistics, the enquiry endpoint, the OpenAPI document, the documentation and the two files every crawler asks for. Every endpoint that returns the record itself requires a key or a first-party origin. You can confirm this from outside — a request to `/api/v1/events` without either returns 401. The count is derived from the gate's own list and published at `/api/v1/facts`.

Deny by default, and measured rather than asserted

Both figures below are measured by the running system on every request to `/api/v1/facts`; this page was built against it on 11 October 2026 and the build refuses to publish a figure that has drifted:

- **Row-level security is enabled on 92 of 92 tables** across the analytical schema and the customer plane. On the customer plane it is forced, the policies are the only route to a row, and the connection that serves customers cannot bypass them.
- **The public database role can execute 0 of the 201 database functions.** Default privileges are set so that a new function does not arrive publicly executable, and the count is re-measured on every request to `/api/v1/facts`.

Stated because a reviewer would find it: the connection that serves the public record holds the database owner role, which row-level security does not constrain. Its safety rests on the application — the licence filter in the shared query builders, parameterised statements, and routes that write only to the tables they own — and on the network and TLS restrictions

above. Moving that path to a role of its own, as the customer plane already has, is scheduled for 1 January 2027 and is the one item on this page we would rather were stronger today.

Default privileges are configured so that a newly created function does not arrive publicly executable, which prevents the permission set widening again over time.

Writes to the customer plane run only through `SECURITY DEFINER` functions that check the caller's capability and record an audit event in the same transaction. The public API's own write paths — alert rules, enquiries, the OSINT ingest — each write one table through a parameterised statement; nothing a request carries is executed as SQL.

Four layers of monitoring, running continuously

1. **Service health.** `/api/v1/health` and `/api/v1/status` answer from the running system. Both are public.
2. **Per-source health.** Every scheduled data fetch records its own outcome — records retrieved, records written, latency, and the reason if it failed. Degraded sources are shown publicly.
3. **Scheduled-job health.** Every job's last run is checked against its schedule, so a job that stops silently is detected.
4. **Data integrity checks.** A set of automated checks that test whether work that completed successfully actually accomplished anything — a cursor that reports progress without advancing, a source rewriting unchanged records, a filter matching nothing because a column is empty, or a licence filter that has stopped filtering.

Findings are recorded when a problem appears and again when it clears, with a daily summary of anything still open. Alerts fire on state change rather than on a fixed interval, so the channel stays meaningful.

Delivery, precisely. Alerts are computed and recorded on every run, and a state change is written with the transition that caused it. The open findings are published on [the status page](#) and on `/api/v1/health` as `integrity`. Outbound delivery goes to the operations mailbox through the platform's email binding; the destination address was verified on 12 September 2026. A transition and the daily digest are sent there; a warning never pages. This paragraph is here because "delivered to an on-call channel" was stated before it was true, and the date it became true is recorded.

What is logged

- Every request made with an issued API key — path, status, row count and timing — retained for 12 months for billing, abuse handling and incident reconstruction.

- Every source run, every scheduled job, and a full revision history for every record, so any figure can be traced back to the fetch that produced it.
- Platform request logs at the edge, retained under the provider's policy.

We do not log what an anonymous visitor viewed. There is no session identity to attach it to.

A server-side request log exists and is filtered per key. It is produced by us on request — during an investigation, or as the audit export named on access control.

Separately, an **append-only audit record** of privileged actions — role changes, key issue and revocation, member invitation and acceptance, operator status changes, publication — is written in the workspace database. It records the actor, the organisation, the action, the prior and new state and a reason. It is append-only by database grant rather than by convention: no role holds UPDATE or DELETE on it, including ours.

There is not yet a customer-facing console to read either one in. The record exists; the screen does not. **Role-based access does exist** and is enforced by the database on every request. **Single sign-on and SCIM do not exist** — they are not built, not configured and not available, and no part of the workspace should be read as offering them.

A small team, with controls sized accordingly

- **Production access is limited to the engineering owner.** The full list of holders, with grant dates, is available under NDA.
- **An automated operating agent** assists with operation under the owner's identity and credentials, on the owner's endpoint. It holds no credential of its own and is bound by the same policies.
- **Everyone with any access signs a confidentiality undertaking** before receiving it, whether employee or contractor.
- **Background checks** are committed for any role with production access, from the first hire, to the extent applicable employment law permits.
- **Security responsibilities will be set out in the engagement terms** of every future hire or contractor rather than in annual training nobody reads; today the code of conduct and the information security policy in the policy set bind everyone with access.
- **Access is revoked within one business day** of a planned departure and immediately on an involuntary one, with rotation of any credential the departing person could have seen.

A team this size has no risk of forgotten contractor access and does not have the separation of duties a large organisation means by the phrase. Both are described in full under access control and development.

Inherited from the infrastructure providers

Vaethra operates no data centre, no office server and no physical infrastructure. Physical and environmental security is the responsibility of the providers, each of which maintains attestations covering it:

- **Cloudflare** — SOC 2 Type II, ISO 27001, ISO 27701, PCI DSS.
- **Supabase**, running on **AWS** — AWS holds SOC 1/2/3, ISO 27001, ISO 27017, ISO 27018 and PCI DSS covering the facilities.

Development endpoints are full-disk encrypted and screen-locked, and are covered by the access control policy. No production data is stored on them.

Ownership, policies and review

- **The engineering owner is accountable for security.**
- **Written policies** cover access control, change management, incident response, cryptography, acceptable use, vendor management and business continuity. They are provided under NDA rather than published, because they describe operational detail.
- **A risk register** records each risk with its likelihood, impact, owner and treatment. It is reviewed quarterly and after any incident.
- **This documentation is reviewed quarterly** and after any material change. Each page carries the date of its last review.
- **Subprocessors are assessed before engagement** and reviewed annually. See [subprocessors](#).

Where to send what

SUBJECT	ADDRESS
Vulnerability reports and security incidents	security@vaethra.com
Security questionnaires, evidence requests, DPAs	security@vaethra.com
Support and enterprise escalation	support@vaethra.com , and +1 253 402-5005 for enterprise Sev-1
Privacy and data subject requests	privacy@vaethra.com
Contracts, terms and licensing	legal@vaethra.com
General enquiries	contact@vaethra.com
Data providers, about our crawling	ops@vaethra.com


Please do not run automated scans against the API without arranging it first. Scanner traffic is indistinguishable from abuse and will be rate-limited and logged as such. Write to us and we will agree a window.

Architecture

How the system is put together: every component, every trust boundary, every external service, and what moves between them.

TC 02Vaethra reference

Four planes, and what crosses between them



Three properties of this design answer most of the questions that follow.

- **Connections to upstream sources are outbound only.** No provider connects to us. We fetch on a schedule, and the request carries nothing beyond a Vaethra user-agent string.
- **A request never waits on an upstream provider.** Reads are answered from our own store, so a provider outage affects freshness rather than availability, and is shown on the public sources panel.

- **The customer plane's connection cannot bypass row-level security.** It connects as a role with no table privileges of its own, and reaches a customer's rows only by a transaction-scoped `SET ROLE` after the organisation is known. The public record is served on a second connection that holds the owner role; that record is public, and the database refuses both connections from anywhere but Cloudflare's address ranges and without TLS. The credentials live in Hyperdrive's configuration, not in the application's secrets.

Production, preview and local, and what each can reach

ENVIRONMENT	PURPOSE	REACHES PRODUCTION DATA
Production	api.vaethra.com, terminal.vaethra.com, atropos.vaethra.com, vaethra.com	Yes
Preview	Per-branch deployments of the static site	No — static HTML only, no secrets bound, no database binding
Local development	The engineer's machine	Yes, for the public record: there is no second copy of the data to develop against, so the machine holds the Worker's secrets in a git-ignored file and reaches the production database through the same pooled connection. That is treated as production access — one holder, full-disk encryption, the access-control page applies. The test suites run against fixtures and never touch production.

There is no shared staging environment holding a copy of production data. This is deliberate: a staging copy is a second copy of the data under weaker controls, and is a common source of exposure. Changes are validated by an automated suite against fixtures, then deployed as an immutable version that can be rolled back in about a minute. See [development](#).

Every flow, and what it contains

FLOW	DIRECTION	CONTENTS
Source ingestion	Vaethra to provider, outbound	An HTTP request and a user-agent. Where a provider requires a key, ours. Nothing relating to any customer.
Record write	Scheduled job to PostgreSQL, through Hyperdrive	Public-source records: events, entities, observations, graph edges, and a revision hash for each
API read	Client to application to cache or PostgreSQL	The record, filtered by licence class. The request is logged against the key.
Terminal and Atropos read	Browser to application	The same records. No account, no cookie, no session identifier.
Text extraction	Application to Workers AI, inside Cloudflare	Headline and summary text from public sources
Embeddings	Application to Workers AI, inside Cloudflare	Event titles and entity names
Bulk extract	Scheduled job to object storage to client	Pre-built event files, excluding restricted licence classes
Webhook alert	Application to the URL you registered	The event that matched your rule
Operational alert	Application to the status page and the operations mailbox	Problem descriptions and counts. No customer data.

Inference runs inside Cloudflare and text does not leave it

- **Two models, both on Cloudflare Workers AI**, which executes within Cloudflare's own network: `@cf/meta/llama-3.1-8b-instruct-fast` for text extraction and relationship assessment, and `@cf/baai/bge-base-en-v1.5` for embeddings.
- **No external model provider is called.** There is no request to OpenAI, Anthropic or any other model API on any code path.
- **What is sent:** headline and summary text from public sources, and entity names.
- **What is never sent:** customer data of any kind.
- **Nothing is used for training.** Inference runs against hosted open-weight models. There is no fine-tuning and no training corpus built from data we hold.
- **Model output is labelled as model output** wherever it appears. An extracted relationship carries its evidence tier and remains distinguishable from a recorded fact.

Every outbound dependency, and whether it sits on the request path

SERVICE	USED FOR	ON THE REQUEST PATH
Cloudflare — Workers, Pages, KV, R2, Durable Objects, Workers AI, Hyperdrive, Access, Turnstile, Web Analytics	The edge platform	Yes
Supabase — PostgreSQL and database functions	The durable record	Yes, on a cache miss
55 public data providers	The record itself	No — scheduled and outbound only
GitHub	Source control and continuous integration	No, and never holds production data

Web fonts are served from our own origin rather than from a third-party font service, so no visitor's IP address is disclosed to one. The Content-Security-Policy on every page of this site, the API, the Terminal and Atropos enumerates the permitted origins and the browser blocks anything else; you can read it in the response headers of this page.

The current position, and what the enterprise workspace adds

Vaethra's record is a single public-source dataset, identical for every reader, and the product holds no customer content. There is therefore no cross-tenant exposure risk: your operational data is not in the system.

The **enterprise workspace** at `workspace.vaethra.com` is the product that changes this, and it is now partly built. It introduces authentication, role-based access and per-customer data isolation. What exists is deliberately narrower than that list suggests: private assets can be held, published in a redacted form and withdrawn; a customer can keep a private dependency map of their own facilities and suppliers, import and export it, and trace a public disruption through it; roles are enforced, and privileged actions are recorded. Member invitation and acceptance exist and are audited; bulk data exports, alerts and integrations in the workspace are **not** built, and single sign-on and SCIM are **not** built.

Isolation is enforced by the database, not by application code: customer rows live in a schema whose tables all carry forced row-level security, reached through a role that cannot bypass it, and a missing predicate yields no rows rather than another organisation's. The negative cases — a session naming a foreign organisation, an unscoped session, a suspended organisation, a key referencing another organisation's service account — are tested rather than asserted. **No customer data is held today.**

The customer records that do exist today — sign-in addresses, API key digests, usage and contact details — are stored in the customer-plane schema of the same database, under

forced row-level security, and reachable only through a role that cannot bypass it.

What is reachable from where

TARGET	REACHABLE FROM
Public API routes	The internet, with a key or a first-party origin. Rate-limited per caller.
Administrative routes	Identities permitted by the identity-aware proxy, and then only with the application bearer token
PostgreSQL	The application through Hyperdrive, from Cloudflare's published address ranges, over TLS. The database refuses any other source and any plaintext connection (measured 11 October 2026). The provider's own management path is the one other route, behind that account's credentials.
Object storage	The application. Extracts are served through the API, not from a public bucket URL.

Cloudflare's DDoS protection and its managed web-application-firewall ruleset sit in front of all of it, and the API applies its own per-caller rate limits behind them. Bot management beyond that ruleset is deliberately not enabled: on this plan it would challenge legitimate API clients as readily as abusive ones.

Access control

Multi-factor authentication, least privilege, production access, joiners and leavers, credential handling and service accounts.

TC 03Vaethra reference

Required on every account, with no exception process

- **Multi-factor authentication is the standing requirement on every administrative account** — Cloudflare, Supabase, GitHub, the domain registrar and the mail provider. The enrolment state of each account is recorded in the access-control evidence and checked at every access review; an account found without a second factor is enrolled the same day. No account with production reach may be entered with a password alone.
- **The factor is an authenticator app on every platform** — time-based one-time codes. SMS is not used as a factor anywhere.
- **Recovery codes are held on paper and in an end-to-end-encrypted drive**, separate from the password manager, and a second holder is named in the access review.
- **There is no MFA exception process.**

Four principals, and what each can do

PRINCIPAL	CAN	CANNOT
Anonymous reader (Terminal, Atropos)	Read the public record through the applications	Execute any database function, reach any administrative route, or obtain a licence-restricted record
API key holder	Read the record within the scope and rate limit of the key	Write anything, reach administrative routes, or obtain records excluded by the licence filter
The application	Serve the public record on the owner-role connection; serve a customer's rows only through a transaction-scoped role that cannot bypass row-level security	Read another organisation's rows on the customer plane, or bypass a policy there
Engineering owner	Deploy, run migrations, and read and change production configuration and data	Recover a stored secret or an issued API key in plaintext — neither is retrievable by anyone

The database enforces its own half rather than relying on the application: row-level security is enabled on 92 of 92 tables, and the public role has execute permission on 0 of the 201 functions. Both measured 11 October 2026 from `/api/v1/facts`.

One holder, with grant dates available on request

- **Production access is held by the engineering owner and by nobody else.** The list is supplied under NDA with grant dates. An automated operating agent assists with operation under the owner's identity and credentials, on the owner's endpoint; it holds no credential of its own and is not a second holder.
- **Administrative endpoints are gated twice** — an identity-aware proxy at the edge, then a bearer token inside the application. Compromising one does not grant access.
- **There is no standing shell access.** No SSH, no bastion, no jump host, because there is no server. Administrative action is an authenticated HTTP call or a platform console operation, both logged by the platform.
- **No shared accounts.** Every human principal is individually identified; automation uses service credentials not attached to a person.
- **The development machine is production access.** It holds the Worker's secrets in a git-ignored file and reaches the production database, because there is no second copy of the data to develop against; it is held by the one production access holder, full-disk encrypted, and covered by this page. The test suites use fixtures.

Quarterly, and on any event that changes who should have what

- **Quarterly**, every account on every platform is reviewed against the question of whether that principal still needs that level of access. The first review was recorded on 12 September 2026; the next falls in December 2026.
- **Immediately** after a departure, a role change, or an incident.
- **API keys are reviewed alongside the account they belong to.** A key with no request in 90 days is flagged as dormant in the workspace and revoked at the next review unless its owner claims it.
- **Each review is recorded** with the date, the reviewer and what changed. These records form part of the evidence pack.

Both directions, with stated deadlines

Vaethra has one engineer and has had no joiner and no leaver. The procedure below is the committed one; it applies from the first, and each step is recorded when it is taken.

Onboarding

1. Confidentiality undertaking signed before any access is granted.
2. Background check where the role carries production access, to the extent applicable employment law permits.
3. Individual accounts created with MFA enforced from first login.
4. Access granted at the minimum level for the role. Production access is not granted by default and is not granted during onboarding.
5. Security responsibilities set out in the engagement terms.

Offboarding

1. Within one business day of a planned departure, and immediately on an involuntary one, every account is disabled.
2. Any shared credential the person could have seen is rotated, not merely revoked from them.
3. Device access is removed and company data removed from personal devices.
4. A written record is kept of what was revoked and when.

Where each kind of secret is held

CREDENTIAL	STORED	RECOVERABLE
Application secrets — provider keys, shared secrets	Platform secret store, encrypted, injected at runtime	No. A lost secret is rotated rather than recovered.
Customer API keys	SHA-256 digest in the database	No. Shown once at creation.
Database credentials	Cloudflare Hyperdrive's configuration and the platform secret store. Never in the repository.	No. Rotated, never recovered.
Human account passwords	A password manager with MFA on the manager itself	By the holder only
MFA recovery codes	Encrypted offline store	By the holder only

- **No secret is present in the repository**, in a tracked configuration file, or in build output. Every push is scanned for credential patterns in the pipeline, and the scan is the first gate the deploy waits on.
- **Rotation is annual as a minimum**, and immediate on any suspicion of exposure, on a departure, or on a provider's advice.

- **Customer key revocation clears the edge cache immediately**, so a revoked key stops working on the next request.

Three, each scoped to a single function

ACCOUNT	SCOPE
Application to database	Two pooled connections held by Cloudflare Hyperdrive: the owner role for the public record, and a role that cannot bypass row-level security for the customer plane.
Deployment	A scoped platform token held by the CI system, able to deploy the Worker and the Pages projects and nothing else. It cannot read data. Issued 12 September 2026.
Provider API keys	Read-only credentials issued by upstream data providers, each scoped to that provider's own data

None is attached to a person, so none breaks when a person leaves. All three are in the rotation schedule and each is separately revocable.

How your side is identified, limited and logged

Everything above is about our access to the system. This section is about yours, because a security review asks both and the page previously answered only one.

Identity

Access to the API is by **API key**. Keys are issued and members are administered in the workspace, which has accounts — email sign-in with a second factor, and roles enforced by the database. The API itself has none: an integration carries a key, not a session, so there is no account takeover, no session fixation, no password reset flow and no credential-stuffing surface on the data path. When someone leaves your team you remove their workspace membership and, if they held one, rotate the key.

Prefix	Every issued secret starts <code>vtx_</code> , so a leaked key is greppable in a repository scan and obvious in a log.
Storage	Compared as a SHA-256 digest. The secret is displayed once, at issue, and is not recoverable afterwards — including by us.
Scopes	Every key carries the read scopes it was issued with and nothing else; the registry of scopes is shown beside the issue form.
Rate limit	Per key, per minute, at the platform rate of 600. Enforced by the platform's rate limiter before the request handler runs.
Row cap	Per request. A larger <code>limit</code> is clamped to 5,000 rather than refused, and <code>total</code> always reports the true match count, so a page never poses as the whole. A lower cap recorded on a key is not yet enforced.
Expiry	Optional per key. A trial key is issued with one.
Revocation	Immediate, and the key stops working on the next request. There is no cache to wait out.
Last used	Recorded per key, and a key unused for 90 days is shown as dormant.

Segregation, without seats

Issue one key per system rather than one per person — the ingestion job, the dashboard, the model, the notebook. Each gets its own rate limit, its own row cap and its own audit trail, so one team's backfill cannot exhaust another team's quota, and revoking one does not interrupt the others. That is the same isolation a role model gives you, expressed as keys, and it does not require us to hold a directory of your staff.

Audit

Every request made with an issued key is recorded: the key, the path, the query, the status, the row count, the bytes and the elapsed milliseconds. A per-key usage summary is shown in the workspace; the full request log for a key is produced on request. Denied anonymous requests are recorded too, deduplicated, against a salted digest of the requesting address rather than the address itself.

What we do not have, and what would change it

Stated plainly because a questionnaire will ask, and because a vendor who lists a control they do not operate is worth less than one who does not.

CONTROL	STATUS	WHAT IT DEPENDS ON
SSO / SAML	NOT BUILT	Requires user accounts, which the product deliberately does not have. If your procurement requires SSO, say so on the first call — it changes the shape of the product, not the timeline of a feature.
SCIM provisioning	NOT BUILT	Same dependency.
Per-seat roles	NOT BUILT	Same dependency. Per-key segregation above is what exists.
IP allow-listing per key	AVAILABLE	An exact-match list per key, checked on every request; a key with a list and an unidentifiable caller is refused.
Customer-visible audit export	PLANNED	The per-key log exists and is produced on request; the export screen in the workspace does not exist yet.

Data protection

Classification, encryption, retention, deletion, residency, export and ownership across the whole life of the data.

TC 09Vaethra reference

Four classes, and which of your data falls into each

CLASS	CONTENTS	HANDLING
Public record	Events, entities, observations and graph edges derived from 55 public feeds — the great majority of what we hold	Served under the licence of its source. Every record carries its publisher, its licence and both timestamps: when the source published it and when we first saw it.
Licence-restricted record	Records from sources whose terms do not permit republishing them individually	Ingested and used for analysis; never served as an individual record and never included in an extract
Customer business record	API key digests, request logs, contact details, contract and billing records	The only class in which a customer is identifiable. Under forced row-level security on the customer plane and reachable only through a role that cannot bypass it. Not shared with anyone.
Operational data	Source run records, job health, cached responses, integrity check findings	Internal, summarised publicly on status

The customer-content class exists in the workspace — a customer's own facilities, routes and assets — and is the most protected class here: forced row-level security, a connection that cannot bypass it, and the negative cases tested rather than asserted. No customer content is held in it today.

Three classes of source, enforced in code and in the database

CLASS	SOURCES	SERVED AS RECORDS	IN A BULK EXTRACT
Open	46	Yes	Yes
Non-commercial	7	Yes, with the licence attached to every record	No
Internal	2	No — neither events nor observations	No

Figures measured 11 October 2026 and readable from `/api/v1/licences`, which reports the same numbers from the running system alongside the count of restricted records held and the count publicly visible, which is zero.

How it is enforced:

- The class is a required field on every source definition, so a new source cannot be added until its licence has been read and recorded.
- The filter is implemented in the shared query builder rather than at each endpoint, so a new endpoint cannot omit it.
- The database enforces it independently, with its own flag and its own condition, so an application error does not defeat it.
- An automated check probes for leakage and alerts if it finds any, and the code and the database are compared against each other on every operations report.

Algorithms by state

STATE	PROTECTION
In transit, public	TLS 1.2 minimum, TLS 1.3 preferred. HSTS for one year including subdomains. HTTP is redirected and never served.
In transit, internal	TLS through Cloudflare Hyperdrive. The database refuses plaintext connections and any source outside Cloudflare's address ranges.
At rest, database	AES-256, including backups and replicas
At rest, object storage	AES-256
At rest, secrets	Encrypted in the platform secret store, not readable back
API keys	SHA-256 digest. The key is never stored, so a copy of the database yields no usable credential.
Endpoint devices	Full-disk encryption, screen lock

Key management is performed by the platforms. We do not operate a key hierarchy of our own. Customer-managed keys are not offered.

How long each class is held, and why

DATA	RETENTION	REASON
The public record — events, entities, observations, graph	Indefinite	It is the product. A disruption archive that forgets cannot support the correlation and backtesting the system exists for.
Record revision history	Indefinite	Provenance. Any figure must be traceable to the fetch that produced it, including a figure that later changed.
API request logs	12 months	Billing, abuse handling, incident reconstruction
Source run records	Rolling window	Operational health
Cached responses	Seconds to hours	Performance. Not a system of record.
Database backups	7 days standard, 28 days on enterprise	Recovery. See business continuity .
Webhook alert rules	Until deleted, or 120 days after the last delivery attempt or save	So that an unowned rule does not continue firing
Contract and billing records	7 years	Statutory tax and accounting retention

What can be deleted and how quickly

- **API key records:** revoked immediately on request; the record is deleted at the end of the contract, subject to statutory retention on the billing record.
- **Webhook alert rules:** deleted immediately on request, or by you at any time from the Terminal.
- **Contact details:** deleted on request, subject to the same statutory retention.
- **Request logs:** expire on the 12-month schedule and can be purged earlier on request.
- **Deletion completes within 30 days** of a verified request, across primary storage and caches. Data in backups is removed as those backups age out of the retention window; backups remain encrypted and access-controlled throughout.

The public record is outside the scope of a deletion request: it is not customer data and it comes from a published source. Where a source corrects or withdraws a record, our copy carries the correction with its revision history intact.

Media disposal is performed by the infrastructure providers, whose attestations cover secure disposal. See [subprocessors](#).

Where data is processed

DATA	LOCATION
The analytical database and its backups	AWS us-east-2, United States (standard); EU region offered on enterprise agreements, provisioned at contract
Cached responses, extracts, coordination state	Cloudflare's global network, cached at the location nearest the reader
Model inference	Cloudflare's network; text does not leave it
Source control	United States. Never holds production data.

Vaethra Technologies LLC is established in the United States, so for a customer in the EEA or the UK we are the data *importer*: transfers are covered by the 2021 Standard Contractual Clauses, incorporated in the [DPA](#) together with the supplementary measures set out there. Both infrastructure providers maintain their own transfer mechanisms.

EU data residency is offered on enterprise agreements. The database platform supports EU regions; where residency is a requirement, it is agreed at contract, a project in an EU region is provisioned within one business day of the order, and the record is migrated to it before the go-live date the order states. It is a provisioning step, not a setting, and the agreement says so.

Who owns what, and how you obtain a copy

- **You own your data.** Anything you provide — contact details, alert rules, and in the enterprise workspace your own facilities and routes — remains yours. We claim no licence over it beyond operating the service, and we do not use it to train anything.
- **The public record is licensed rather than owned.** It comes from public sources under their own terms, and what you may do with it is set by those terms, which travel with every record in a `license` field. We cannot grant rights the upstream publisher did not grant us.
- **Export is a standard feature.** The API returns JSON and CSV, bulk extracts are pre-built files, and the OpenAPI document is public. There is no export fee and no exit fee.
- **On termination,** a complete export of anything you supplied is provided within 30 days on request, in a machine-readable format.
- **We do not sell, share, rent or broker data.** There is no advertising, no analytics broker and no data broker anywhere in this business.

Privacy

What personal data Vaethra processes, why, where, for how long, and the rights individuals have over it.

TC 10Vaethra reference

The design holds very little personal data

The Terminal and Atropos require no sign-in. There is no account, no profile, no record of what any visitor viewed, and no cookies of any kind — not analytics cookies, not advertising cookies, not session cookies. That is why there is no consent banner. Watchlists and layouts are held in your own browser's storage and never reach a server. The one place a person signs in is the workspace, where an API customer's staff administer keys and members; it holds their sign-in addresses and nothing about anyone else.

Web fonts are served from our own origin, so viewing the site discloses your IP address to nobody but us and our CDN. There are no third-party trackers, no advertising network, no analytics broker and no social widgets on any Vaethra property. The public products — this site, the Terminal and Atropos — carry Cloudflare's cookieless page counts, placed on the page by us; the customer workspace and the operator console carry no third-party script.

This is an architectural property rather than a policy one, and it is the main reason a compromise of Vaethra would expose very little about any individual.

This page is the detailed explanation for enterprise review. The formal notice is the [Privacy Policy](#), and the processor terms are the [DPA](#).

Every category, with its basis and retention

CATEGORY	WHOSE	PURPOSE	LAWFUL BASIS	RETENTION
Business contact details — name, work email, organisation	Customers and enquirers	To respond, to operate the contract, to send required notices	Contract, or legitimate interests for an enquiry	Duration of the relationship, then deleted on request
Enquiries — name, email, organisation, message, the page it was sent from, and a salted digest of the sending address	Enquirers	To reply, and to tell a person's submission from an automated one	Legitimate interests	Until answered and closed; deleted on request
Workspace accounts — sign-in email, session and second-factor records, role	An API customer's staff	Sign-in and authorisation in the workspace	Contract	Until the membership is removed
Denied requests — a salted digest of the address, the path and the status	Any caller the API refused	Abuse handling	Legitimate interests	12 months
API key records — key digest, owner, scope, usage	API customers	Authentication, rate limiting, billing, abuse handling	Contract	Contract duration; logs 12 months
Request logs — path, status, timing, row count, key	API customers	Billing, abuse handling, incident reconstruction	Contract and legitimate interests	12 months
Alert rule destinations — the webhook URL you register	Anyone who creates one	To deliver the alert while your browser is closed	Legitimate interests, at your request	Until deleted, or 120 days after the last delivery attempt or save
Edge request metadata — IP, user-agent, timestamp	Every visitor	Security, rate limiting, abuse prevention	Legitimate interests	Per Cloudflare's retention policy
Aggregate page counts	Site visitors	Understanding which pages are used. Cookieless, measures no individual, and disclosed here.	Legitimate interests	Aggregate only

Not processed at all: special-category data, children's data, biometric data, location data about individuals, browsing history, and any data obtained from a data broker. There is no advertising and no profiling.

Where the dataset itself names an individual

Vaethra's record is drawn from public sources about physical disruption, and a small proportion of it names people — a sanctioned individual on a published government list, a named official in a public advisory, an author credited in a public dataset. This is public-interest processing of information already published by an authority or a publisher, and every record carries the source and licence it came from.

- We do not enrich, cross-reference or build profiles of individuals.
- We do not add anything the source did not publish.
- Where a source corrects or withdraws a record, our copy carries the correction with its revision history intact.
- An individual who believes a record about them should not be held should write to privacy@vaethra.com. We assess it against the public-interest and freedom-of-expression grounds and respond with a reason either way within one month.

Controller or processor, and where

- **For business contact details and API key records, Vaethra is the controller.** The [Privacy Policy](#) is the notice.
- **For anything you supply for us to process on your behalf** — today, alert rule destinations; in the enterprise workspace, your own facilities and routes — **Vaethra is the processor** and the [DPA](#) governs it.
- **Processing locations:** the database and its backups are in AWS us-east-2, with an EU region offered on enterprise agreements and provisioned at contract. Edge processing is global, at the location nearest the reader.
- **Transfers** outside the EEA rely on the 2021 Standard Contractual Clauses, incorporated in the DPA with supplementary measures.

What can be requested, and the response time

RIGHT

WHAT WE DO

Access	Provide a copy of everything held about you, in a machine-readable format
Rectification	Correct anything inaccurate. For a public-source record we correct our copy and retain the revision history.
Erasure	Delete it, subject to statutory retention on billing records and to the public-interest grounds above
Restriction	Stop processing while a dispute is resolved
Portability	Machine-readable export of anything you supplied
Objection	Stop legitimate-interests processing unless we can demonstrate overriding grounds, which we will state
Complaint	To us, and to the supervisory authority where you live or work regardless. Vaethra is established in the United States and has no lead authority of its own.

We respond within one month, extendable by two months for a complex request, in which case we tell you within the first month and give the reason. There is no charge. Write to privacy@vaethra.com.

There is no automated decision-making with legal or similarly significant effects on any individual. Vaethra assesses ports, routes and infrastructure, not people.

Ownership and records

- **The founder is accountable for privacy.** Vaethra is not required to appoint a data protection officer under Article 37: its core activities do not involve large-scale systematic monitoring of individuals or special-category data.
- **A record of processing activities** is maintained under Article 30 and is available under NDA.
- **Privacy by design is structural.** The public products have no accounts, no cookies and no profiles because they were built that way; there is no setting that turns tracking on.
- **Data protection impact assessments** are carried out where processing is likely to be high risk. The enterprise workspace is assessed before the first customer's own data is held in it; it holds none today.
- **Every subprocessor is assessed for privacy** — the current five on 12 September 2026, any new one before engagement — and reviewed annually.

Subprocessors

Every third party that stores or processes data on Vaethra's behalf, what they hold, where, and their own attestations.

TC 11Vaethra reference

Five, and no others

SUBPROCESSOR	PURPOSE	DATA	REGION	ATTESTATIONS
Cloudflare, Inc. United States	Edge compute, CDN, DNS, web application firewall, object storage, key-value store, access control, model inference	Request metadata, cached API responses, bulk extracts, cached satellite imagery, public-source text sent for extraction and embedding	Global edge	SOC 2 Type II, ISO 27001, ISO 27701, PCI DSS
Supabase, Inc. United States	Managed PostgreSQL — the durable record — and database functions	The event and entity record, API key digests, alert rules	AWS us-east-2, or an EU region on enterprise agreements	SOC 2 Type II
Amazon Web Services, Inc. United States	Underlying infrastructure for Supabase. Named explicitly rather than left behind our direct vendor.	As Supabase — they host it	us-east-2, Ohio	SOC 1/2/3, ISO 27001, ISO 27017, ISO 27018, PCI DSS
GitHub, Inc. United States	Source control and continuous integration, which holds the deployment token and the migration token as encrypted secrets	Source code and deployment credentials. No production data, no customer data.	United States	SOC 1/2 Type II, ISO 27001
Proton AG Switzerland	The mailboxes behind every @vaethra.com address: support, security, privacy, legal, contact and operations	Whatever a customer or enquirer writes to us, and our replies	Switzerland	ISO 27001

Data providers are not subprocessors. The 55 public feeds we read receive nothing from us but an HTTP request and a user-agent string. They are listed separately under coverage, with the licence of each.

Web fonts are self-hosted, so no font service receives any visitor's IP address.

There is very little customer personal data to pass on

The public products have no user accounts, so no name, email or profile is attached to use of the Terminal or Atropos, and nothing about a session reaches a subprocessor beyond the request metadata any web service necessarily sees.

Three exceptions, stated precisely:

- **Webhook alert rules.** The destination URL you register is stored in the database so the rule can be evaluated while your browser is closed. It is owned by an anonymous token your browser generates, not by an identity.
- **API key and workspace records.** The digest, the owner, the usage and the sign-in addresses of a customer's workspace members are stored for billing, abuse handling and authorisation. They live in Supabase on AWS.
- **Correspondence.** An email to any of our addresses, and an enquiry once it is forwarded to the operations mailbox, is held by Proton.

No advertising network. No analytics broker. No data broker. No sale or sharing of any data with anyone outside the table above.

Assessed, and reviewed annually

The five subprocessors in use were assessed on 12 September 2026, after engagement; the assessments are in the evidence pack. Any addition is assessed before engagement, in the order below.

1. **Necessity.** Whether the requirement can be met without adding a third party. The list is five long because the answer is usually yes.
2. **Security assessment.** Their attestations, security documentation, incident history and their own subprocessors.
3. **Privacy assessment.** What personal data they would process, the transfer mechanism, and their retention and deletion terms.
4. **Contract.** A data processing agreement incorporating the Standard Contractual Clauses where a transfer requires them, with terms no weaker than those we give you.
5. **Annual review**, and immediate review after any incident of theirs that affects us.

30 days' notice, with a right to terminate

We give 30 days' notice by email to enterprise customers before adding or replacing a subprocessor, and this page carries the current list with the date it was last reviewed.

If a change is unacceptable to you, that is grounds to terminate without penalty for the remainder of the term, with a full export of anything you supplied. That right is in the DPA.

To join the notification list, write to support@vaethra.com.

Availability

Uptime as measured, the monitoring behind it, maintenance windows, and the enterprise service level.

TC 04Vaethra reference

Published on an endpoint anyone can read

Availability is sampled continuously and published at api.vaethra.com/api/v1/uptime, which needs no key and no account. It reports availability per day, the sample coverage behind each figure, and how many data feeds were reporting.

Sample coverage is published beside the percentage because a figure drawn from four samples is a different claim from one drawn from a thousand, and a status page that shows only the percentage does not let you tell them apart.

Continuous sampling began recently, so the published window is currently short and lengthens daily. Samples are retained for **400 days**, which the endpoint publishes as `retentionDays`, so the series grows into a rolling twelve months and then holds there. If your procurement process requires twelve months of history today, we do not have it — ask for the date the window reaches your requirement and we will give you the arithmetic rather than an estimate.

This sentence used to say "we do not have it yet" against a 90-day retention, which meant twelve months was never going to arrive. The retention was raised rather than the sentence softened.

What has to fail before service is affected

- **The edge is stateless and globally distributed.** The API, Terminal and Atropos run in Cloudflare's network at hundreds of locations. There is no primary region to lose, and a failed location is routed around by the platform without action from us.
- **A read never waits on a data provider.** Provider calls run on a schedule, off the request path. An upstream outage affects freshness and is shown on the sources panel; it does not affect availability.
- **One data source cannot take the system down.** Each of the 55 sources runs in isolation with its own time budget and records its own health.

- **Responses are cached at the edge by class**, so a traffic spike is absorbed by the CDN rather than by the database. The most expensive responses are additionally pre-built on a schedule.
- **The database is the one component without a hot standby in the current design.** It is managed, backed up and monitored, and recovery is covered under business continuity. While it is unreachable, cached responses continue to serve.
- **Authorisation fails closed, and so does caching.** A refused request never reads from or writes to the shared cache, so a credential problem cannot become a data exposure.

Four layers

LAYER	WATCHES	PUBLIC
Service health	<code>/api/v1/health</code> and <code>/api/v1/status</code> , answered by the running system	Yes
Source health	Every source's last run: records retrieved, records written, latency, failure reason	Yes, on the sources panel
Scheduled-job health	Every job's last run against its schedule	No — job names and error text are operational
Data integrity checks	Whether work that completed successfully accomplished anything	Summarised on <u>status</u>

The fourth layer is the one worth explaining. Conventional monitoring answers whether a job ran. These checks answer whether it did what it was for: a cursor incrementing its counters without advancing, a source rewriting unchanged records, a filter matching nothing because the column it reads is empty, a licence filter that has stopped filtering.

Findings are recorded when a problem appears and again when it clears, with a daily summary of anything still open, and the open set is published on the status page. Alerts fire on state change rather than on a fixed interval, so that a channel people watch stays worth watching. Transitions and the daily digest are delivered to the operations mailbox; the security overview records when that became true.

Routine deployments cause no downtime

- **Deployments are zero-downtime.** A new version is uploaded and traffic moves to it; there is no restart and no connection drain.
- **Deployments run from the pipeline on every change to the default branch**, at whatever hour the change lands. The engineering owner is the one who lands it and is therefore present for it, and the pipeline's post-deploy verification reports what shipped.

- **Schema changes are additive first.** A column is added, written, backfilled and only then read from, so the old and new application versions can run at the same time.
- **Where impact is expected, enterprise customers receive 5 business days' notice** by email with the window, the expected effect and the rollback plan, and never inside a customer's stated business hours without agreement.
- **Emergency maintenance** — a security fix that cannot wait — is announced as it begins and is followed by the same written report an incident receives.

What is committed, by tier

	FREE (TERMINAL, ATROPOS)	TEAMS	ENTERPRISE
Availability commitment	None	None	99.9% monthly , with service credits
Measured against	The <code>availability</code> field of the public uptime endpoint anyone can read — not <code>uptime</code> . See below.		
Support response (Sev-1)	Best effort	4 business hours	1 hour, 24/7
Status	<u>Public, including failures</u>		

The Terminal and Atropos are free and offered as they are, without an availability commitment.

The enterprise commitment is 99.9% monthly availability of the API, measured against the `availability` field of the public uptime endpoint you can read yourself, with service credits against the following month's fee. The following are excluded from the calculation:

- Scheduled maintenance announced 5 business days in advance.
- Upstream data provider outages. A provider outage is a degraded source, visible on the sources panel, and is not unavailability of Vaethra. This distinction is written into the agreement rather than left to interpretation.
- Force majeure, and faults in the customer's own network.

The credit schedule is published in the Enterprise Service Addendum, which an enterprise order form incorporates.

Business continuity

Backups, retention, recovery objectives, and the procedure for each disaster scenario.

TC 05Vaethra reference

Four independent copies of different things

ASSET	BACKED UP	WHERE	RETENTION
The analytical database — events, entities, observations, graph, API key digests	Daily automated snapshots; continuous (point-in-time) recovery is switched on for the project at the start of an enterprise agreement	Encrypted, stored separately from the primary	7 days standard, 28 days on enterprise agreements
Bulk extracts and the satellite-scene cache	Written to object storage as produced	Cloudflare R2 — a different provider from the database	Rolling
Raw source payloads	Not yet archived. A content-addressed archive of every fetched payload, held with a second provider, is a written requirement and not a built system.	—	—
Schema and infrastructure definition	Every change, as a versioned migration in source control	GitHub, and every engineer's local clone	Permanent, full history
Application code and configuration	Every commit; each deployment is an immutable version	GitHub and the platform's version history	Permanent

Backups are encrypted at rest and held separately from the primary, so a compromise of the running database is not a compromise of its history.

The record can also be rebuilt from source material: everything in the database derives from public feeds we can fetch again, and the schema and the code that fetch them are in source control with a second provider. Total loss of the database would be slow and expensive to recover from, but it would not be permanent loss of the product. What it would lose is any history a source has since withdrawn, which is why the raw archive above is a requirement.

Committed objectives, by tier

OBJECTIVE	STANDARD	ENTERPRISE
RPO — maximum data loss	24 hours	5 minutes
RTO — time to restore service	4 hours	2 hours
Backup retention	7 days	28 days
Recovery granularity	Daily snapshot	Any point in the retention window

Point-in-time recovery, which is what brings the recovery point objective to minutes, is a configuration change on the database platform: it is switched on for the project at the start of an enterprise agreement, not built. These are the objectives we commit to contractually.

The edge does not require recovery in the same sense. The API, Terminal and Atropos are stateless and globally distributed: a failed location is routed around automatically, and a faulty deployment is recovered by redeploying the previous immutable version, which takes about a minute.

Five scenarios and the response to each

SCENARIO	IMPACT	RESPONSE	EXPECTED
Faulty deployment	API errors or incorrect output	Redeploy the previous immutable version. No restore, no data involved.	Under 5 minutes
Edge location or region failure	None visible	The platform routes around it. No action required.	Automatic
Database unavailable	Cached reads continue; fresh reads and writes fail	Platform failover, or restore to a new instance and repoint the connection layer	Within RTO
Database corrupted or destroyed	Loss of the durable record	Restore to the recovery point, repoint, verify with the automated integrity checks, then re-ingest any gap from the public sources	Within RTO; any gap backfills behind it
Provider account compromised	Potentially total	Incident response: revoke every credential, rotate every secret, restore from backup into a clean project, rebuild the edge from source control. Everything required is in the repository.	Same day to service

The published summary of the runbook

The full runbook is operational and is provided under NDA. Its structure:

1. **Declare.** Classify the incident and start the notification deadlines in incident response. Nothing waits for the technical work to finish.
2. **Contain.** Stop the harm before diagnosing it — disable the affected path, roll back, or revoke the credential.
3. **Restore.** Recover the database to a new instance from the chosen recovery point, repoint the connection layer, and redeploy the edge from the last known good version.
4. **Verify before declaring recovery.** Run the automated integrity checks against the restored system and read the output. A restored database that has lost an index, a trigger or a licence filter looks healthy until it is queried.
5. **Backfill** the window between the recovery point and the failure from the public sources.
6. **Report.** A written post-incident report within 5 business days.

What is tested and how often

- **Rollback is the same mechanism as an ordinary deployment** — redeploying an earlier immutable version — and a deliberate rollback to a prior version is part of the restore test, so that it is exercised rather than assumed.
- **Restore from backup is tested to a scratch instance**, at least annually and before the first enterprise agreement, and the result written up: what was restored, from which backup, how long it took, and what failed. The test plan exists; the first exercise has not yet been run, and its report joins the evidence pack when it has. Testing that a backup restores is a different thing from testing that it exists.
- **The automated integrity checks are the verification step**, which is what allows a restore to be trusted: they establish that the recovered system is correct rather than merely running.
- **A full disaster-recovery exercise runs annually** against the written test plan, the first before the first enterprise agreement, with the report shared with enterprise customers on request.

The key-person question

Vaethra is a small company, and the material continuity risk is people rather than infrastructure.

- **Nothing depends on anyone's memory.** Schema, configuration, site content and deployment are all in source control, and the repository carries the reasoning behind each

decision as well as the change itself. That documentation is a continuity control.

- **Credential recovery is not to depend on one person being reachable.** Recovery codes for every platform exist today, on paper and in an end-to-end-encrypted drive, and a second holder is named in the access review. The key-person risk stays in the risk register, because a second holder is a mitigation and not a second engineer.
- **Enterprise agreements can include source-code escrow** through an escrow agent engaged for that agreement, releasing on defined insolvency or cessation events.
- **The data is reproducible.** Every record derives from a public source that still exists.

Incident response

How an incident is detected, classified, contained and communicated, with the deadlines we commit to in writing.

TC 06Vaethra reference

Three levels, defined by impact

LEVEL	DEFINITION	ACKNOWLEDGED	UPDATES
Sev-1	Service unavailable; data accessed without authorisation; or the record is incorrect in a way a customer could act on	1 hour, 24/7	Hourly until contained, then every 4 hours until resolved
Sev-2	Major degradation — a class of feed down, materially stale data, or severe performance loss	4 business hours	Daily
Sev-3	Minor or cosmetic; a single source degraded with the rest unaffected	2 business days	On resolution

A single failing data source is normally Sev-3 by design: sources run in isolation so that one provider cannot take the system down, and the failure is shown publicly on the sources panel while everything else continues.

Any suspected unauthorised access is treated as Sev-1 from the first minute, before it is confirmed, so that notification deadlines are not missed while it is being investigated.

The Sev-1 line is answered by the engineering owner, with a named second responder behind it and no rotation beyond the two; [support](#) says so. Severity is set by impact rather than by who reported it. If you believe something has been graded too low, tell us and it will be regraded.

Four sources of detection

1. **Automated monitoring**, described under [availability](#): health checks, per-source run records, scheduled-job health, and the data integrity checks. State changes are recorded with a daily summary of anything still open.
2. **Customer reports**, which are not treated as lower priority than automated detection. support@vaethra.com, or the enterprise phone line for a Sev-1.

3. **Security researchers** — security@vaethra.com. See [vulnerability management](#).
4. **Provider status**. Cloudflare and Supabase publish status pages and notify the account holder of incidents affecting the services we use.

Six steps

1. **Declare and classify**, within the acknowledgement window, stating what is known so far including where the cause is not yet known. Declaring starts every deadline on this page.
2. **Contain before diagnosing**. Disable the affected path, roll back the deployment, or revoke the credential. Diagnosis follows containment.
3. **Notify**. Affected enterprise customers within 24 hours of classification, by email and, for a Sev-1, by phone. Where personal data may be involved, the 72-hour requirement below applies in parallel.
4. **Investigate**. Establish the cause, the scope and the time window. Preserve logs and evidence before changing anything that would remove them.
5. **Remediate and verify against the running system** rather than against a test.
6. **Report**. A written post-incident report within 5 business days: what happened, the cause, what was affected, what was done, and what has changed to prevent recurrence.

Notification within 72 hours

Where an incident involves unauthorised access to systems or data, affected customers and any relevant supervisory authority are notified within 72 hours of Vaethra becoming aware, in line with GDPR Article 33. Vaethra is established in the United States and has no lead authority of its own: as a processor, we notify the customer, and the customer notifies the authority with jurisdiction over them.

The notice covers, per Article 33(3), the nature of the breach, the categories and approximate number of data subjects and records concerned, the contact point, the likely consequences, and the measures taken or proposed.

The personal-data surface is small: workspace sign-in addresses, API key records, business contact details and enquiries. We will nonetheless notify on any confirmed unauthorised access to production systems, whether or not personal data was involved, so that customers can make their own assessment of the risk.

Where a customer has its own regulatory reporting obligation — NIS2, DORA, or a sector regulator — the 24-hour customer notice is set to leave time inside their deadline. If your deadline is shorter, we will agree it in the contract.

Hourly updates on a Sev-1, with or without news

- During a Sev-1 you receive an update every hour, including when the content is that work is continuing.
- Enterprise customers receive updates from a named person rather than a queue address.
- Service failures appear on status and material incidents are added to the changelog.
- An incident is not described as resolved until it has been verified resolved against the running system.

Every report ends with a change

A post-incident report is required to close with a concrete change: a code change, a new automated check, or a documented decision. In this system it is usually the second.

The automated integrity checks described under availability are the accumulated output of this process. Each one exists because a specific failure was found that had reported success, and each now detects that failure the next time it occurs.

Reports are provided to affected customers, and the record of them forms part of the evidence pack available under NDA.

Vulnerability management

Scanning, dependency monitoring, patch windows, security testing and responsible disclosure.

TC 07Vaethra reference

Continuous, against a deliberately small surface

- **Automated dependency scanning** runs against the repository: vulnerability alerts and security-update pull requests are enabled on it, and a weekly grouped pull request refreshes the rest, so a vulnerable package arrives as a pull request rather than an email that goes unread.
- **Every install is pinned by a lockfile**, so builds are reproducible and a compromised upstream release cannot enter on the next install.
- **The dependency surface is small by design**. The API is built on a minimal edge framework and the platform's own primitives. The marketing site has no build-time JavaScript dependencies and is generated by a script in the repository.
- **No third-party JavaScript executes on our pages** beyond Cloudflare's own analytics and challenge scripts, both enumerated in a restrictive Content-Security-Policy.
- **The runtime is patched by the platform**. There is no operating system, container image or interpreter for us to leave unpatched.

By severity, timed from when we become aware

SEVERITY	DEFINITION	REMEDIED WITHIN
Critical (CVSS 9.0–10.0)	Remotely exploitable without authentication, in a component we reach	24 hours , as emergency maintenance if required
High (7.0–8.9)	Serious and exploitable in our configuration	7 days
Medium (4.0–6.9)	Exploitable given preconditions we partly meet	30 days
Low (0.1–3.9)	Limited impact, or not reachable in our configuration	90 days , or accepted with a written rationale and a date

Where a finding is not exploitable in our configuration, the assessment is documented rather than the rating being reduced silently.

What runs before every deployment

- **2,834 automated tests** and static type checking run on every push, and the pipeline that runs them is the deployment path: the deploy job exists only on a run in which the scan, the type check, the four test suites and the builds have passed.
- **Security-relevant behaviour has its own tests.** The licence filter, the authentication gate, the rate limiter, the key-digest comparison and the alerting logic each have tests asserting that they refuse, not only that they permit.
- **The database's posture is measured rather than assumed.** Row-level security on 92 of 92 tables and the public role's 0 executable functions are queried from the running database on every request to `/api/v1/facts`.
- **Automated integrity checks run continuously in production** and detect the class of failure a scanner cannot: a filter that has stopped filtering, a cursor reporting progress it did not make, a column empty on every row so a security condition matches nothing.
- **Every push is scanned for credential patterns** before anything else in the pipeline runs.

An independent penetration test is to be commissioned before the first enterprise agreement, and the report will be available under NDA once complete. Internal security review is continuous and has produced findings, one of which is described below.

September 2026: database function permissions

Internal review found that a set of database functions were executable by the platform's public role. Reaching them still required a credential that has only ever existed in server-side secrets, and no data was exposed, but the permission was wider than intended and the review treated it as a defect rather than an accepted risk.

Resolution. Every database call was routed through a secret-gated function running inside the database environment with a fixed operation allowlist. That was deployed and every affected endpoint verified. Execute permission was then revoked from the public role on all functions, and default privileges were changed so that a newly created function cannot arrive publicly executable. Measured afterwards: the public role can execute none of the 77 functions, and the database returns a permission error on any attempt.

The same work surfaced an unrelated defect — an endpoint that had been returning an empty result rather than an error when a call failed — which was also fixed.

Since then the design has moved on: on 9 September 2026 the secret-gated function was retired and the database is reached through two pooled connections, one of which cannot bypass row-level security (see [architecture](#)). The public role's execute set is now measured on every request and published at `/api/v1/facts` : 0 of 201 functions on 11 October 2026.

We publish findings of this kind because a reviewer is entitled to see how we handle one.

How to report, and what you receive

Report to security@vaethra.com with enough detail to reproduce.

STAGE	COMMITMENT
Acknowledgement	2 business days
Initial assessment and severity	5 business days
Progress updates	Every 7 days until closed
Fix	Per the patch windows above
Credit	In the changelog if you want it

What we ask

- **Do not run automated scans against the API without arranging it first.** Scanner traffic is indistinguishable from abuse and will be rate-limited and logged as such. Write to us and we will agree a window.
- Do not access, modify or destroy data that is not yours, and do not degrade the service for others.
- Allow reasonable time to fix before publishing. We will not ask for indefinite silence, and we will not take legal action against anyone reporting in good faith under these terms.

We do not currently run a paid bug bounty. Reports are acknowledged, fixed and credited.

Our contact details are also published at [/.well-known/security.txt](#) in the format defined by RFC 9116.

Secure development

How a code change reaches production, what gates it, and the controls around production change.

TC 08Vaethra reference

Six steps, in order

1. **Change, with its reasoning in the commit.** With one engineer, a change lands on the default branch; every push runs the pipeline below, and a branch-and-review step is the committed procedure from the day a second engineer joins.
2. **Static type checking.** The codebase is fully typed and a type error stops the change. This is what makes controls such as the licence classification a compile-time requirement rather than a convention.
3. **Automated tests.** 2,834 of them, including tests that assert security controls refuse.
4. **Review** against the recorded reasoning in the repository before the change lands — today by the one engineer who wrote it; by a second engineer from the day one joins.
5. **Deploy as an immutable version.** The platform records every deployment; nothing is edited in place.
6. **Verify against the running system** after deployment, rather than treating a passing test suite as confirmation that a feature works.

Version control is the audit trail

- **Every change lands through version control** with a full auditable history: who, when, what and why. The reasoning is recorded in the commit alongside the change.
- **Database changes are versioned migrations**, applied in order, recorded in a ledger and replayable from the repository. There are currently 154. An ad-hoc change to the production schema that is not written as a migration is treated as a defect, and the operations report compares what the code expects against what the database holds.
- **Deployments are immutable and versioned.** Rollback is redeploying a prior version rather than editing a running system.
- **Schema changes are additive first**, so the old and new application versions can run concurrently during a deployment.

- **Secrets are never committed.** Every push is scanned for credential patterns before anything else in the pipeline runs.

What each environment can reach

ENVIRONMENT	DATA	SECRETS
Local	The production database, through the same pooled connection: there is no second copy of the record. Test runs use fixtures.	The Worker's secrets, in a git-ignored file. The machine is treated as production access.
Preview	Static HTML only	None bound
Production	The live record	Platform secret store, injected at runtime, not readable back

There is no shared staging environment holding a copy of production data. A staging copy is a second copy of the data under weaker controls and is a common source of exposure. Correctness is established by the type system, the test suite and the integrity checks; deployment safety is established by immutability and fast rollback.

The development machine reaches production and is treated as production access under access control: one holder, full-disk encrypted, its secrets in a git-ignored file that the credential scanner would refuse if it were ever staged.

The specific controls in this codebase

- **Parameterised queries throughout.** No string-built SQL on any path that accepts user input. Query filters are built by shared builders so that a new endpoint cannot construct its own and lose a control — the licence filter is implemented inside one of those builders for that reason.
- **Fail closed.** Authorisation failures deny; a refused request never reads from or writes to the shared cache; an unrecognised data source is treated as licence-restricted rather than open.
- **Inputs are validated and bounded.** Every list endpoint has a hard row ceiling, every request has a time budget, and every external fetch has a timeout. Bulk export is a licensing arrangement rather than a query parameter.
- **Output encoding and a restrictive Content-Security-Policy** on every public page, enumerating the permitted origins.
- **No secret reaches a log, an error message or a response body.**
- **Errors are surfaced rather than swallowed.** A component that cannot do its job reports that, rather than returning an empty result — an empty result and a broken component are

indistinguishable downstream. This is the most consistently enforced rule in the codebase.

What governs a change at this company's size

Reviewers ask whether one person can put arbitrary code into production unnoticed. The current position:

Deployment is performed by the engineering owner. A second human approver is planned as the team grows and will be documented here when it is in place. The controls that apply today:

- **Deployment runs from the pipeline, after the gates.** The deploy job exists only on a run in which the credential scan, the type check and the 2,834 tests — including tests that assert the security controls refuse — have passed, and it proves what it shipped by reading the deployed commit back from the running system. The engineering owner also holds a deploy credential for emergency use; any use of it is recorded in the [changelog](#).
- **Every deployment is recorded and immutable.** There is a complete, externally held record of every version that has run and when, and rollback takes about a minute.
- **The database enforces its own half.** Deny-by-default row-level security and zero publicly executable functions mean an application error does not become a data exposure.
- **The integrity checks monitor outcomes, not intent.** A change that breaks the licence filter, empties a column or stops a cursor advancing is detected regardless of who deployed it.
- **Every change is recorded with its reasoning,** so the history is reviewable after the fact, including by an auditor.

Compliance

Framework status, and a control-by-control map of SOC 2, ISO/IEC 27001 Annex A and the GDPR against where each control is implemented.

TC 12Vaethra reference

Where Vaethra stands, framework by framework

FRAMEWORK	STATUS	DETAIL
SOC 2 (Security, Availability, Confidentiality)	BUILT TO THE CRITERIA AUDIT NOT STARTED – ON THE ROADMAP	Vaethra is building the application and its operating procedures around the AICPA Trust Services Criteria, and the control map in §2 sets out each criterion against where it is implemented. The audit itself has not been started and no firm is engaged. The plan: a Type I report on the Security criteria, with the firm engaged in the first quarter of 2027, and the Type II observation window after it; this row will name the engagement when one exists. Until then the audit rights in the DPA stand in its place.
ISO/IEC 27001:2022	BUILT TO ANNEX A CERTIFICATION NOT STARTED – ON THE ROADMAP	The information security management system is being built to the Annex A control set, mapped in §3. Certification by an accredited body has not been started. Vaethra is not certified and does not claim to be; this row will change when a certification body is engaged.
GDPR / UK GDPR	OBLIGATIONS APPLY	Not a certification, and not a status anyone has conferred: the GDPR is a legal obligation, and whether an organisation meets it is a conclusion for a regulator or a court, not a badge a controller awards itself. No independent assessment, DPIA or external legal review has been carried out. Vaethra Technologies LLC is established in the United States and processes personal data falling under the GDPR as a processor for its customers. A Data Processing Agreement is available at /legal/dpa , incorporating the 2021 Standard Contractual Clauses and the UK Addendum. Article-by-article mapping in §4.
CCPA / CPRA	NOT IN SCOPE	Below the Act's revenue and consumer thresholds, and Vaethra does not sell or share personal information as those terms are defined. There is no advertising, no data broker and no cross-context behavioural profiling.
PCI DSS	OUT OF SCOPE	Vaethra never receives, transmits or stores cardholder data. Invoicing is by bank transfer, and any future card payment would be handled entirely by a PCI-compliant processor.
HIPAA	OUT OF SCOPE	Vaethra does not process protected health information and does not sign Business Associate Agreements. The public epidemiological feeds we read carry no individual health data.
NIS2	NOT IN SCOPE	Vaethra is not established in the EU and is not an in-scope essential or important entity under any national transposition. Where a customer is, our notification timelines are set to support their reporting obligations — see incident response .

FRAMEWORK	STATUS	DETAIL
EU AI Act	NOT ASSESSED	No legal assessment has been made of whether Vaethra is a provider of an AI system under the Act. Language models are used for text extraction, relationship assessment and embedding. No biometric processing, no automated decisions about individuals, no emotion recognition. Model output is labelled as model output wherever it appears. See architecture §4 .
TISAX, FedRAMP, IRAP, C5	NOT HELD	Not held and not currently in progress. If one is a requirement for your organisation, raise it early and we will give you a straight answer on scope and timing.

Each Trust Services Criterion, and where it is implemented

Read against the 2017 Trust Services Criteria (2022 revision). Every row links to the page that describes the control in full, so this table can be used as an index rather than as a summary.

CC1 — Control environment

CRITERION	HOW IT IS MET	WHERE
CC1.1 Integrity and ethical values	Written code of conduct in the policy set; a confidentiality undertaking is required before any access is granted, from the first person after the founder	Security §7
CC1.2 Board oversight	The engineering owner is accountable for security; the risk register is reviewed quarterly and after any incident	Security §9
CC1.3 Structures and reporting lines	Security ownership, escalation path and named contacts documented	Security §9–10
CC1.4 Competence	Background checks for any role with production access from the first hire; security responsibilities set out in engagement terms	Access control §5
CC1.5 Accountability	Individual accounts only, no shared credentials; every change attributable in version control	Development §2

CC2 — Communication and information

CRITERION	HOW IT IS MET	WHERE
CC2.1 Quality information	Health, status and uptime answered by the running system on public endpoints; per-source health published	<u>Status</u>
CC2.2 Internal communication	Findings recorded on state change, with a daily summary of anything open	<u>Availability §3</u>
CC2.3 External communication	This trust centre; incident communication commitments; 30 days' subprocessor notice; 90 days' API deprecation notice	<u>Trust centre, Support §5</u>

CC3 — Risk assessment

CRITERION	HOW IT IS MET	WHERE
CC3.1 Objectives specified	Security objectives documented in the information security policy	Policy set, under NDA
CC3.2 Risks identified and analysed	Risk register: risk, likelihood, impact, owner, treatment	Register, under NDA
CC3.3 Fraud risk considered	Assessed as part of the register; separation of read and write paths; no payment handling in scope	Register, under NDA
CC3.4 Significant change assessed	Risk reviewed on architecture change and after every incident	<u>Security §9</u>

CC4 — Monitoring of controls

CRITERION	HOW IT IS MET	WHERE
CC4.1 Ongoing evaluations	Four layers of monitoring, including automated integrity checks that test whether controls are still working rather than whether jobs completed	<u>Availability §3</u>
CC4.2 Deficiencies communicated	Alerts computed on state change and recorded; findings summarised publicly on the status page	<u>Status</u>

CC5 — Control activities

CRITERION	HOW IT IS MET	WHERE
CC5.1 Control selection	Controls selected against the risk register and implemented in code where possible rather than in policy	Development §4
CC5.2 Technology controls	Deny-by-default database authorisation, licence filtering in a shared query builder, fail-closed authorisation and caching	Security §5 , Data protection §2
CC5.3 Policies and procedures	Documented policy set covering access, change, incident, cryptography, vendors and continuity	Policy set, under NDA

CC6 — Logical and physical access

CRITERION	HOW IT IS MET	WHERE
CC6.1 Logical access security	MFA the standing requirement on every administrative account, enrolment checked at each access review, no exception process; API keys stored as SHA-256 digests; row-level security on 92 of 92 tables, forced on the customer plane	Access control §1–2
CC6.2 Registration and authorisation	Access granted at minimum level for the role; production access not granted by default; keys issued individually with scope	Access control §3
CC6.3 Access modification and removal	Quarterly access reviews; revocation within one business day of a departure and immediately on an involuntary one, with credential rotation	Access control §4–5
CC6.4 Physical access	No Vaethra-operated facilities. Inherited from Cloudflare and AWS.	Security §8
CC6.5 Asset disposal	No physical media held. Disposal is performed by the infrastructure providers under their own attestations.	Data protection §5
CC6.6 External threat protection	DDoS protection and the managed WAF ruleset at the edge; per-caller rate limiting; the database listener restricted to Cloudflare's address ranges and to TLS	Architecture §7
CC6.7 Transmission and movement of data	TLS 1.2 minimum with HSTS; the database reachable only from the edge, over TLS; extracts served through the API rather than from a public bucket	Data protection §3
CC6.8 Malicious software prevention	No general-purpose compute to infect; immutable deployments; pinned dependencies with continuous scanning; commit scanning for credentials	Vulnerability management §1

CC7 — System operations

CRITERION	HOW IT IS MET	WHERE
CC7.1 Vulnerability detection	Continuous dependency scanning; configuration review; automated integrity checks in production	<u>Vulnerability management §1, §3</u>
CC7.2 Monitoring for anomalies	Health, source, job and data integrity monitoring; rate-limit and authorisation failure logging	<u>Availability §3</u>
CC7.3 Security event evaluation	Three-level severity model with defined acknowledgement windows	<u>Incident response §1</u>
CC7.4 Incident response	Six-step process: declare, contain, notify, investigate, remediate, report	<u>Incident response §3</u>
CC7.5 Recovery from incidents	Immutable rollback in about a minute; restore procedure with integrity verification before recovery is declared	<u>Business continuity §4</u>

CC8 — Change management

CRITERION	HOW IT IS MET	WHERE
CC8.1 Change authorisation and testing	The pipeline gates: credential scan, type check, 2,834 automated tests, builds, then deployment from the pipeline with a post-deploy proof of what shipped. Database changes are versioned migrations recorded in a ledger.	<u>Development §1-2</u>

CC9 — Risk mitigation

CRITERION	HOW IT IS MET	WHERE
CC9.1 Business disruption mitigation	Stateless global edge, isolated data sources, encrypted backups, committed recovery objectives, annual DR exercise	<u>Business continuity</u>
CC9.2 Vendor and partner management	Subprocessors assessed before engagement and reviewed annually; contractual terms no weaker than those we give customers; 30 days' notice of change	<u>Subprocessors §3-4</u>

A1 — Availability, C1 — Confidentiality

CRITERION	HOW IT IS MET	WHERE
A1.1 Capacity management	Database size and per-table growth measured on every operations report, on a platform whose disk scales automatically; edge caching absorbs load spikes	Security §6 , Availability §2
A1.2 Backup and recovery infrastructure	Encrypted daily backups stored separately from the primary; schema and code with a second provider; a raw-payload archive is a stated requirement not yet built	Business continuity §1
A1.3 Recovery testing	Rollback is the deployment mechanism and is exercised deliberately in the restore test; a restore test and an annual DR exercise committed against a written plan, the first run before the first enterprise agreement	Business continuity §5
C1.1 Confidential information identified	Four-class data classification with defined handling for each	Data protection §1
C1.2 Confidential information disposed of	Deletion within 30 days of a verified request across primary storage and caches; backups age out within the retention window	Data protection §5

The 2022 control set, by theme

Annex A of ISO/IEC 27001:2022 groups 93 controls into four themes. The rows below cover the controls applicable to Vaethra's scope; controls relating to facilities we do not operate are marked as inherited from the infrastructure providers, each of which publishes its own attestations — listed per provider on [/security/subprocessors](#).

A.5 — Organisational controls

CONTROL	HOW IT IS MET	WHERE
A.5.1 Policies for information security	Documented policy set, reviewed at least annually and after any incident	Policy set, under NDA
A.5.2 Information security roles	The engineering owner is accountable; responsibilities set out in engagement terms	Security §9
A.5.7 Threat intelligence	Provider advisories, dependency advisories and CISA KEV monitored; CISA KEV is itself one of our ingested sources	Vulnerability management §1
A.5.8 Security in project management	Security review is part of the change process rather than a separate gate	Development §1
A.5.9–5.11 Asset inventory, acceptable use, return of assets	Asset register maintained; acceptable use in the policy set; offboarding removes device access and company data	Access control §5
A.5.12–5.14 Classification, labelling, transfer	Four-class classification; every record carries its source and licence; transfer restricted by class	Data protection §1, §2
A.5.15–5.18 Access control, identity, authentication, rights	Least privilege, individual identities, MFA everywhere, quarterly review	Access control
A.5.19–5.22 Supplier relationships	Subprocessors assessed before engagement, contractually bound, reviewed annually, and published	Subprocessors
A.5.23 Cloud services security	Two providers, both attested; configuration reviewed; no self-managed infrastructure	Security §2
A.5.24–5.28 Incident management	Documented process with severity model, notification deadlines, evidence preservation and post-incident report	Incident response
A.5.29–5.30 Continuity and ICT readiness	Committed recovery objectives, five documented scenarios, a written restore test plan	Business continuity
A.5.31–5.34 Legal, IP, records, privacy	Licence classification enforced in code and database; GDPR programme; statutory records retention	Data protection, Privacy
A.5.35–5.37 Review, compliance, procedures	Quarterly review of this documentation and of access; control-by-control mapping maintained here	This page

A.6 — People controls

CONTROL	HOW IT IS MET	WHERE
A.6.1 Screening	Background checks for any role with production access from the first hire, to the extent applicable employment law permits	Access control §5
A.6.2 Terms and conditions of employment	Confidentiality undertaking required before access is granted, from the first person after the founder	Security §7
A.6.3 Awareness and training	Security responsibilities in engagement terms; practice documented in the repository alongside the code	Development §4
A.6.5 Responsibilities after termination	Access revoked within one business day, credentials rotated, record kept	Access control §5
A.6.7 Remote working	Full-disk encrypted, screen-locked endpoints; no production data stored locally	Security §8
A.6.8 Reporting security events	Published reporting address and disclosure process with commitments	Vulnerability management §5

A.7 — Physical controls

CONTROL	HOW IT IS MET	WHERE
A.7.1–7.14 Physical perimeter, entry, equipment, media, disposal	Inherited. Vaethra operates no data centre, office server or physical infrastructure. Cloudflare and AWS hold ISO 27001, ISO 27017 and SOC 2 attestations covering these controls.	Security §8
A.7.7 Clear desk and clear screen	Endpoint screen lock enforced; no printed material	Security §8
A.7.9 Security of assets off-premises	Full-disk encryption on all endpoints; access through the same authenticated paths as any other client	Security §8

A.8 — Technological controls

CONTROL	HOW IT IS MET	WHERE
A.8.1 User endpoint devices	Encrypted, screen-locked, no production data at rest	Security §8
A.8.2–8.5 Privileged access, information access, source code, secure authentication	One production access holder; deny-by-default database authorisation; MFA everywhere; source access individually identified	Access control
A.8.6 Capacity management	Database size and per-table growth measured on every operations report	Security §6
A.8.7 Protection against malware	No general-purpose compute; immutable deployments; pinned dependencies	Vulnerability management §1
A.8.8 Management of technical vulnerabilities	Continuous scanning with severity-based patch windows: 24 hours critical, 7 days high, 30 days medium, 90 days low	Vulnerability management §2
A.8.9 Configuration management	All configuration in version control; deployments immutable and versioned; database schema as ordered migrations	Development §2
A.8.10 Information deletion	Deletion within 30 days of a verified request; backups age out	Data protection §5
A.8.11 Data masking	Test runs use synthetic fixtures; the development machine reaches production and is treated as production access; API keys stored only as digests	Development §3
A.8.12 Data leakage prevention	Licence filter enforced in the shared query builder and independently in the database, with an automated leakage probe	Data protection §2
A.8.13 Information backup	Daily encrypted snapshots stored separately, point-in-time recovery switched on for enterprise agreements; restore test committed	Business continuity §1
A.8.14 Redundancy	Stateless globally distributed edge; automatic routing around a failed location	Availability §2
A.8.15–8.16 Logging and monitoring	API request logging retained 12 months; source, job and integrity monitoring; alerting on state change	Security §6
A.8.17 Clock synchronisation	Platform-managed time on all compute; every record carries both the source's timestamp and ours	Data protection §1
A.8.19–8.20 Software on operational systems, network security	No operating system to manage; DDoS protection and the managed WAF ruleset at the edge; the database listener restricted to the edge's address ranges	Architecture §7

CONTROL	HOW IT IS MET	WHERE
A.8.21 Security of network services	All service-to-service traffic over TLS; the database enforces TLS and a source allow-list; the customer plane reached only through definer functions	Architecture §1
A.8.22 Segregation of networks	Production and preview separated; preview has no production binding; the development machine is treated as production access	Architecture §2
A.8.23 Web filtering	Restrictive Content-Security-Policy on every public page — this site, the API, the Terminal and Atropos; the customer workspace and the operator console carry frame, content-type, referrer and permissions policies, and a Content-Security-Policy for them is committed	Architecture §5
A.8.24 Use of cryptography	TLS 1.2+ in transit, AES-256 at rest, SHA-256 for credential digests; documented cryptography policy	Data protection §3
A.8.25–8.29 Secure development lifecycle	Secure coding standard, review, 2,834 automated tests including tests that assert controls refuse, static type checking, environment separation	Development §1, §4
A.8.30 Outsourced development	Not applicable — development is not outsourced	—
A.8.31 Separation of environments	No shared staging environment holding production data	Development §3
A.8.32 Change management	Pipeline gates, deployment from the pipeline, immutable versioned deployments, rollback in about a minute	Development §1–2
A.8.33 Test information	Synthetic fixtures only; production data is never copied into a test environment	Development §3
A.8.34 Protection during audit testing	Audit access arranged in advance with scope agreed; scanning against production coordinated rather than ad hoc	DPA §9

Where each obligation is discharged

ARTICLE	OBLIGATION	HOW IT IS MET
Art. 5	Principles	Data minimisation is architectural — no accounts, cookies or profiles on the public products. Purpose limitation and storage limitation documented per category. Privacy §2
Art. 6	Lawfulness	A lawful basis stated for every category of processing. Privacy §2
Art. 12–22	Data subject rights	Access, rectification, erasure, restriction, portability and objection, answered within one month. Privacy §5
Art. 24, 32	Controller responsibility, security	Technical and organisational measures documented in full. DPA §14
Art. 25	Data protection by design and default	No accounts, no cookies and no trackers on the public products; self-hosted fonts. There is no setting that turns tracking on. Privacy §1
Art. 28	Processor obligations	Executable DPA with subprocessor terms, audit rights and deletion obligations. DPA
Art. 30	Records of processing	Maintained; available under NDA. Privacy §6
Art. 33–34	Breach notification	72 hours to the supervisory authority and to affected customers, with the Article 33(3) content. Incident response §4
Art. 35	Impact assessments	Carried out where processing is likely to be high risk; the enterprise workspace is assessed before the first customer's own data is held in it. Privacy §6
Art. 27	Representative in the Union and in the UK	Not appointed. Vaethra is established in the United States and, for the visitors and enquirers whose data it holds as a controller, relies on the Article 27(2) exemption: the processing is occasional, involves no special-category data, and is unlikely to result in a risk to individuals — there are no accounts, cookies or profiles on the public products, and an enquiry is held until it is answered. The same position is taken under the UK GDPR. Recorded on 12 September 2026 as the owner's election, to be confirmed by counsel; if it changes, a representative's details will appear here and in the Privacy Notice .
Art. 37	Data protection officer	Not required: core activities do not involve large-scale systematic monitoring of individuals or special-category data. The founder is accountable. Privacy §6
Art. 44–49	International transfers	2021 Standard Contractual Clauses and the UK Addendum, with supplementary measures and government-access commitments. DPA §10–11

Four things, on request

1. **This documentation**, as a single PDF — eighteen pages covering what a standard vendor security review asks, including the control maps above, with the cross-references intact.
2. **Your questionnaire, completed in writing**. CAIQ, SIG Lite or a bespoke spreadsheet.
Committed turnaround: five business days.
3. **The evidence pack, under NDA** — policies, risk register, records of processing, access-control evidence, vendor assessments, and the recovery test report once the first exercise has run. Listed in full, with what exists today, on [the trust centre index](#).
4. **A call with the engineer who built the system**, who can answer architecture questions directly.

If a specific certification or artefact is a condition of purchase, raise it on the first call and we will give you a date or a straight answer that we cannot meet it. [Tell us what you need](#).

Support

Channels, hours, response targets by severity, escalation, and how much notice you get before anything changes.

TC 13Vaethra reference

How to reach us, by tier

TIER	CHANNELS	HOURS
Enterprise	support@vaethra.com and +1 253 402-5005	24 hours a day, 7 days a week for a Sev-1. Business hours (08:00–18:00 CET, Monday to Friday) otherwise, with email monitored outside them.
Teams	support@vaethra.com	Business hours, 08:00–18:00 CET, Monday to Friday
Free (Terminal, Atropos)	support@vaethra.com	Best effort
Security reports	security@vaethra.com	Any tier, any hour. See disclosure .

The phone line is for enterprise customers and for a Sev-1. Hours are given in Central European Time because the engineering owner is based in Vienna; the company is a Wyoming limited liability company. For anything that is not urgent, email gets a better answer, because whoever picks it up can look at the system before replying.

By subject

The table above is by tier. This one is by what you are writing about, and reaching the right address first saves a round trip.

SUBJECT	ADDRESS
Faults, outages, anything not working	support@vaethra.com
Vulnerability reports and security incidents	security@vaethra.com
Security questionnaires, evidence requests, DPAs	security@vaethra.com
Privacy and data subject requests	privacy@vaethra.com
Contracts, terms and licensing	legal@vaethra.com
API access, pricing, everything else	contact@vaethra.com
Data providers, about our crawling	ops@vaethra.com

Time to a named person, by severity

These are response targets — a person with the problem, not an autoresponder. Resolution time depends on the fault and is not committed in advance.

SEVERITY	DEFINITION	ENTERPRISE	TEAMS	FREE
Sev-1	Service unavailable, or the record is incorrect in a way you could act on	1 hour, 24/7	4 business hours	Best effort
Sev-2	Major function degraded; a workaround exists	4 business hours	1 business day	Best effort
Sev-3	Minor fault, question, or documentation gap	1 business day	3 business days	Best effort
Sev-4	Feature request or roadmap question	3 business days	5 business days	Best effort

Severity is set by impact rather than by who reported it. If you believe something has been graded too low, say so and it will be regraded.

During an incident, updates follow the schedule on [incident response](#) rather than this table: hourly for a Sev-1, including when the update is that work is continuing.

What lets us answer in one exchange

- **What you asked for** — the endpoint and parameters, or what you clicked. A copy of the request is ideal.
- **What you received** — the status code and body, or a screenshot.

- **What you expected**, and why. Some reports turn out to be a licence filter or a source's publication cadence working as documented, and knowing what you expected distinguishes those quickly.
- **When**, with a timezone. Request logs are keyed by time and API key.
- **Your key's identifier** — the prefix.

Please never send a live credential, ours or anyone else's. If one has already been sent, tell us and it will be revoked immediately.

If the response is not adequate

1. **Reply on the thread and say you are escalating.** This changes who reads it.
2. **Enterprise: call.** +1 253 402-5005. For a Sev-1, call first and email afterwards rather than waiting for a reply.
3. **Named contacts.** Enterprise agreements name an engineering owner and a commercial owner with direct addresses.

Vaethra is a small company, which means you reach the person who built the system within one step. The 24/7 line is the corporate number — it receives calls and SMS and is answered by that person, with a named second responder who holds the recovery codes and can restore from the runbook if the first is unreachable. There is no follow-the-sun organisation behind it, and an enterprise agreement says so rather than implying one.

How much warning you get before something changes

CHANGE	NOTICE
Breaking API change — a field removed, a type changed, an endpoint retired	90 days by email to every active key holder, and in the changelog . A <code>Deprecation</code> response header on the affected endpoint is planned and not yet implemented, so do not build a client that waits for one.
Additive API change — a new field or endpoint	None. Within v1, fields are added and never removed or repurposed, as the API documentation states; a client should tolerate a field it has not seen.
New or replaced subprocessor	30 days by email, with a right to terminate without penalty if the change is unacceptable. See subprocessors .
Planned maintenance with expected impact	5 business days, and never inside a customer's stated business hours without agreement
Source removed from coverage	30 days where the removal is our decision. Where an upstream provider withdraws a feed, as soon as we know, with the reason recorded in the changelog .
Price change	Never within a committed term. 60 days before a renewal.