Privacy Notice

Effective date: 2026-04-30 · Version 1.0

This notice describes how Value Driven AI ("we") handles personal data through the C2MD Compliance Agent service ("the Service"). It is intended for individuals whose personal data may be processed by the Service, the customers (organisations) using the Service, and prospective customers evaluating the Service.

Who we are

Value Driven AI is a single-operator commercial entity. We operate the C2MD Compliance Agent as a software-as-a-service product hosted in the European Union.

Contact for privacy enquiries: `privacy@getvda.ai`

What this Service does

The Service generates compliance documentation (assessments, governance artifacts, regulatory mappings) for AI agent systems on behalf of customers. Customers submit descriptions of their AI agents to the Service; the Service returns generated compliance bundles.

Our role under data protection law

For personal data submitted by customers via the Service, we operate as a **processor** under the General Data Protection Regulation (GDPR / Regulation 2016/679). The customer is the **controller**.

This means: the customer determines what personal data (if any) is submitted to the Service and for what purpose; we process that data only as instructed by the customer through their use of the Service.

What personal data may be processed

Customers may include personal data within their submitted agent descriptions. The categories of personal data we may process depend on what each customer chooses to submit, and may include:

- Names of individuals (developers, end users, subjects)
- Email addresses
- Phone numbers
- Locations
- Bank account identifiers
- Date and time references that could indirectly identify individuals

We do not solicit, require, or independently collect personal data from customers. The personal data in scope is only what customers choose to include in their agent descriptions.

How we minimise personal data exposure

Before any submitted data reaches our underlying AI processing systems, we apply automated pseudonymisation. Personal data entities are detected and replaced with placeholder tokens (e.g., `<PERSON_1>`, `<EMAIL_1>`) before any inference, screening, or generation operation. The mapping from placeholder back to the original value is held only in transient memory during the request and is not logged, stored, or returned to the caller.

This pseudonymisation does not displace data subjects' rights under GDPR — the customer remains controller of the underlying data — but it does mean that the personal data is not exposed to AI inference systems or persisted by us beyond the lifetime of an individual request.

Where data is processed

All processing occurs within the European Union, specifically the Google Cloud `europe-west1` region (St. Ghislain, Belgium). EU data residency is enforced at the architectural level: the Service's source code hardcodes the EU region, and a misconfigured

deployment cannot route processing outside the EU without a code change.

Sub-processors

We use a single sub-processor: **Google Cloud Platform**. Google operates the underlying compute, AI inference, content screening, secret storage, and audit logging services within the EU. Google's data processing addendum is at cloud.google.com/terms/data-processing-addendum.

We do not use any other third-party sub-processors. There are no analytics services, marketing pixels, or third-party logging providers in our processing path.

If we change our sub-processor list in future, we will notify customers in advance via the customer-facing changelog (once published) before any new sub-processor begins processing customer data.

Retention

We do not persist customer-submitted personal data. All such data exists only transiently in memory during request processing and is discarded when the response is returned. There is no customer-facing database, storage layer, or backup of submitted content.

Audit logs (managed by Google Cloud) retain timestamps and request metadata for compliance and security purposes:

- Administrative audit logs: 400 days (Google default for `_Required` log sinks)
- Application audit logs: 30 days (Google default for `_Default` log sinks)

These logs do not contain customer-submitted personal data — they record request timestamps, service identifiers, latencies, and similar operational metadata.

Vertex AI processing

Customer data passes through Google Vertex AI for inference. Under Google's standard terms:

- Google does not use customer data to train Google's AI models
- Google may retain customer data transiently (typically 24-30 days) for abuse-monitoring purposes
- Data does not leave the EU at any point

We are working to enable Zero Data Retention (ZDR) configuration for Vertex AI calls, which would eliminate the abuse-monitoring retention. ZDR is currently blocked on upstream SDK feature availability and will be enabled when feasible. See our Security & Privacy Overview document for full detail.

Data subject rights

Because we operate as a processor and do not persist personal data, the practical mechanics of data subject rights (access, erasure, portability, rectification, objection) under GDPR Articles 15-22 are best exercised against the controller — typically the customer organisation that submitted the data.

If you are a data subject who believes your personal data was processed by us through a customer's use of the Service, please contact the customer first. They are the data controller and have direct accountability for your rights.

If you cannot identify or reach the controller, you may contact us at `privacy@getvda.ai`. We will work with you to identify the relevant customer and to provide whatever metadata we have from audit logs (timestamps of any requests that may have processed your data).

We will respond to data subject enquiries within 30 days where we can identify the relevant data, or sooner where we cannot.

Security

The Service implements substantial technical security controls including TLS encryption in transit, encryption at rest, multi-factor authentication via delegated identity providers, narrowly-scoped service accounts, and pseudonymisation before inference. A complete description is in our published Security & Privacy Overview at getvda.ai/security.

Security disclosures should be sent to `security@getvda.ai`.

Changes to this notice

Material changes to this notice will be announced with at least 30 days' notice via the customer-facing changelog (once published) and via the version metadata at the top of this document.

Non-material changes (typo corrections, link updates, contact information updates) may be made without advance notice; the version date will reflect the change.

Complaints

If you believe we are not handling your personal data lawfully, you have the right to lodge a complaint with a supervisory authority. The relevant supervisory authority for our processing operations is the **Belgian Data Protection Authority** (dataprotectionauthority.be), as our processing occurs in Belgium (Google Cloud `europe-west1` region, St. Ghislain).

You may also lodge a complaint with the supervisory authority in your country of residence.

This privacy notice is published as part of C2MD's transparency commitments. Document version 1.0, effective 2026-04-30.