Privacy Policy

GRIP — Guernsey Regulatory Intelligence Platform

Effective Date: 19 March 2026
Version 1.0

> **GRIP is operated as a private platform. This Privacy Policy explains what personal data we collect, how we use it, and your rights in relation to it. By using GRIP, you acknowledge that you have read and understood this policy.**

1. Who We Are

GRIP (Guernsey Regulatory Intelligence Platform) is a private platform providing AI-assisted legal and regulatory research services relating to the Bailiwick of Guernsey. For the purposes of applicable data protection law, the operator is the data controller in respect of personal data collected through GRIP.

Guernsey has its own data protection legislation: the Data Protection (Bailiwick of Guernsey) Law, 2017 (as amended), administered by the Office of the Data Protection Authority (ODPA). This policy has been prepared to comply with that law. Where the UK GDPR or EU GDPR is separately applicable to any processing activity, this policy is also intended to address those requirements.

2. What Personal Data We Collect and Store

GRIP uses a purpose-specific storage architecture. The following sets out exactly what data is held in each system.

2.1 Account and Organisation Data (Supabase/PostgreSQL — EU West, Ireland)

We store the following in our primary relational database:

- User identifiers, email address, role, optional full name, membership status, approval metadata, and last sign-in timestamp.
- Organisation (tenant) profile data, plan tier, token allowance, token usage totals, and billing status.
- Organisation membership workflow data (pending, approved, or rejected status).
- Session authentication state managed through Supabase Auth mechanisms and cookies.

2.2 Query and Workspace Data (Supabase/PostgreSQL — EU West, Ireland)

For product functionality and query history, we store:

- The text of queries you submit and the text of AI-generated answers returned to you.
- Citation references and confidence metadata returned with answers.
- Token usage totals and per-query token breakdown fields.
- Matter and workspace records, including title, description, status, and context fields you provide.
- Clarification question history — the answers you give when GRIP asks follow-up questions to narrow a query.
- Saved searches and their associated filter configurations.

2.3 Legal Corpus and Document Data

GRIP maintains a corpus of public Guernsey legal materials. The following data is stored across multiple systems:

- Document metadata (document ID, title, domain, jurisdiction, dates, status, source URL) — stored in PostgreSQL (Supabase, EU West).
- Searchable text chunks and vector/sparse embeddings of those chunks — stored in Qdrant (London, UK).
- Legal relationship metadata (amendment links, supersession relationships between documents) — stored in Neo4j graph database (Vultr infrastructure).
- Raw ingestion artifacts and corpus archive files — stored in AWS S3 object storage under corpus-specific prefixes.

2.4 Customer-Uploaded Private Documents

If you upload private documents to GRIP (for example, to use within a Matter workspace):

- Uploaded PDF content is text-extracted, chunked, and embedded. The resulting chunks and embeddings are stored in Qdrant (London, UK) with strict tenant-level scoping to prevent cross-organisation exposure.
- Corresponding document metadata is stored in PostgreSQL (Supabase, EU West) and linked to your organisation and, where applicable, your workspace.
- Raw uploaded files may be archived in AWS S3.

Do not upload documents containing confidential client information, legally privileged communications, or the personal data of third parties unless you have

satisfied yourself that doing so is lawful and appropriate. Uploaded documents are processed by AI embedding providers (see Section 5).

2.5 Billing and Payment Data

GRIP uses Paddle as its Merchant of Record. This means that when you purchase a subscription or top-up, Paddle is legally the seller — they process the payment, issue the invoice in their own name, and handle VAT and tax compliance on your behalf. We do not store full payment card numbers or CVV data in GRIP's application database. What we do store:

- Paddle customer ID and Paddle subscription ID.
- Top-up purchase records: bundle type, token amounts, expiry timestamps, and Paddle transaction references.
- Subscription status and plan tier as reported by Paddle webhooks.
- Invoice records provided by Paddle for display in your account.
- Billing status and payment state fields.

All payment instrument data (card numbers, bank details) is held solely by Paddle, who process payments under their own terms and privacy policy. Paddle is registered in the USA (Paddle.com Inc.) with its EU operations managed through Paddle.com Europe Ltd.

2.6 Alerts and Communications Preferences

If you enable corpus change alerts, we store:

- Alert subscription preferences: domains, tiers, and delivery channel selected.
- Destination values: email address, webhook URL, or Slack destination URL depending on your chosen channel.
- Alert delivery outcomes and error logs.
- Corpus changelog events that trigger your alerts.

2.7 Audit and Operational Logs

GRIP maintains append-only audit records for security and compliance purposes. Each audit event contains:

- Tenant ID, user ID, action type, resource reference, relevant metadata, IP address, and timestamp.

Audit records are written to PostgreSQL (Supabase, EU West) and may additionally be exported to AWS S3 for compliance retention purposes. Application logs may also contain limited operational diagnostics, including query snippets in some processing pipelines.

2.8 Data We Do Not Intentionally Collect

GRIP is not designed to collect special category personal data as defined under Guernsey data protection law (including data relating to health, political opinions, religious beliefs, biometric identification, or criminal records). Please do not include such data in your queries or uploaded documents. If you do so inadvertently, that data will be processed only to the extent necessary to respond to your request.

3. How We Use Your Personal Data

We process your personal data for the following purposes:

1. Providing the service: to process your queries, return AI-assisted results, manage workspaces and matters, and operate and maintain the GRIP platform.
2. Account and organisation management: to manage user accounts, authenticate sessions, enforce plan limits, manage organisation memberships, and communicate with you about your account.
3. Billing: to process subscription payments and top-up purchases via Paddle (our Merchant of Record). Paddle handles all payment processing, VAT, and invoicing. We do not store payment card data ourselves.
4. Transactional email: to send account verification, password reset, billing notification, and alert emails via Resend (using Supabase Auth as the trigger).
5. Corpus alerts: to notify you of changes to Guernsey legal materials within the domains and tiers you have subscribed to.
6. Audit and security: to maintain an audit trail of platform actions, detect and respond to abuse, and investigate security incidents.
7. Service improvement: to analyse aggregate usage patterns, identify errors, and improve the quality and coverage of GRIP's corpus and AI responses.
8. Legal compliance: to comply with applicable law, regulatory requirements, or court orders.

4. Legal Basis for Processing

We rely on the following legal bases under the Data Protection (Bailiwick of Guernsey) Law, 2017, and, where applicable, the UK GDPR and EU GDPR:

- Contractual necessity: processing required to provide the GRIP service under our Terms of Use, including storing queries, managing accounts, and processing billing.
- Legitimate interests: audit logging, security monitoring, fraud prevention, and aggregate service improvement analytics — where our interests are not overridden by your data protection rights.
- Legal obligation: retention of billing and financial records for the period required by applicable law.

- Consent: where you have given specific consent, such as for marketing communications or certain analytics cookies. You may withdraw consent at any time without affecting the lawfulness of prior processing.

5. Third-Party Data Processors

GRIP relies on the following third-party processors who may process your personal data on our behalf. We have data processing agreements or equivalent contractual arrangements in place with each.

Provider	Role / Data Processed	Location & Privacy Policy
Supabase (PostgreSQL)	Primary relational database — accounts, queries, workspaces, billing state, alerts, audit logs	EU West, Ireland — supabase.com/privacy
Qdrant	Vector database — document chunks, embeddings, private uploaded documents	UK (London) — qdrant.tech/legal/privacy-policy
Neo4j	Graph database — legal relationship metadata	Vultr-hosted (see below) — neo4j.com/privacy
AWS S3	Object storage — corpus archives, audit exports, raw ingestion artifacts	EU region — aws.amazon.com/privacy
Paddle	Merchant of Record — payment processing, VAT, invoicing	USA / EU (Paddle.com Europe Ltd) — paddle.com/privacy
Resend	Transactional SMTP — account and billing emails	USA — resend.com/legal/privacy-policy
Anthropic	LLM (Claude) — query text and retrieval context sent for answer generation	USA — anthropic.com/privacy
Cohere	Reranking — query text and candidate document chunks	USA/Canada — cohere.com/privacy
Voyage AI	Embeddings — query and document text for vectorisation	USA — voyageai.com/privacy
Vultr	Cloud server infrastructure — graph database hosting	USA — vultr.com/legal/privacy
Fly.io	Backend application hosting	USA — fly.io/legal/privacy-policy
Vercel	Frontend hosting and CDN	USA — vercel.com/legal/privacy-policy

AI subprocessors (Anthropic, Cohere, Voyage AI) receive only the data needed to perform each individual request: query text, retrieved document chunks for context, and candidate chunks for reranking. They do not receive account identifiers, billing data, or uploaded document metadata.

6. Where Your Data is Stored — Data Residency Summary

We have designed GRIP's infrastructure to keep personal and account data within the UK and EU where possible. The following is a complete data residency map:

Data Category	Location	System
Account, org, billing state, audit logs, query history, workspaces	EU West (Ireland)	Supabase / PostgreSQL
Document chunks, embeddings, private uploaded documents	UK (London)	Qdrant
Legal relationship metadata	Vultr-hosted server	Neo4j
Corpus archives, audit exports	AWS S3 (EU region)	Amazon Web Services
Payment card data and instruments	Paddle (USA / EU) — not held by GRIP	Paddle.com Inc. / Paddle.com Europe Ltd
Query text sent for AI answer generation	USA (per-request, not persisted by GRIP)	Anthropic / Cohere / Voyage AI
Transactional emails	USA (delivery only)	Resend

> **Account credentials, billing references, query history, audit logs, and structured personal data are all stored within the EU or UK. The only personal data routinely transmitted to the USA is query text sent to AI providers for answer generation, and your email address sent to Resend for transactional delivery. Do not include personally identifiable information of third parties in your queries.**

7. Tenant Isolation and Access Controls

GRIP is a multi-tenant platform. We implement the following controls to prevent cross-organisation data exposure:

- All data records in PostgreSQL include a tenant identifier. Row-level security policies enforce that users can only access records belonging to their own organisation.
- Vector search in Qdrant applies tenant and public/private filters on every retrieval operation, preventing your uploaded documents from being surfaced in other organisations' query results.
- Graph database queries in Neo4j are scoped to the public legal corpus only; no tenant-specific data is stored in the graph layer.
- Bearer tokens issued by Supabase Auth are scoped per user and validated on every API request.

8. Your Rights

Subject to applicable law, you have the following rights in relation to your personal data held by GRIP:

9. Right of access: to obtain a copy of the personal data we hold about you.

10. Right to rectification: to request correction of inaccurate or incomplete personal data.
11. Right to erasure: to request deletion of your personal data. Note that deletion of your account triggers removal or deactivation of user records. If your organisation has no remaining members, tenant-level cleanup is also triggered, removing associated query history, workspaces, and uploaded documents. Audit records required for legal compliance may be retained in anonymised or archived form.
12. Right to restriction: to request that we restrict processing of your data in certain circumstances.
13. Right to data portability: to receive a copy of your personal data in a structured, machine-readable format.
14. Right to object: to object to processing based on our legitimate interests.
15. Right to withdraw consent: where processing is based on consent, to withdraw it at any time without affecting the lawfulness of prior processing.

To exercise any of these rights, please contact us using the details in Section 13. We will respond within one calendar month. You also have the right to lodge a complaint with the Office of the Data Protection Authority (ODPA) at odpa.gg, or with your local supervisory authority if EU or UK GDPR applies to your use.

9. Data Retention

We retain personal data only for as long as necessary for the purposes described in this policy or as required by law. The following retention schedules apply to each data class:

Data Class	Retention Period	Notes
Account and organisation data	Duration of account + 7 years	Post-closure retention for legal and audit purposes
Query history, workspaces, matters, saved searches	Until deleted by user or account closure	No automatic expiry currently implemented
Billing records and payment references	7 years	Financial record-keeping obligation
Top-up token bundles	Per expiry timestamp on each bundle	Explicit expiry set at time of purchase
Audit log records	7 years (PostgreSQL); may be archived to S3 for compliance	Append-only; not deletable by users
Application and operational logs	Up to 90 days	Rotated automatically
Private uploaded documents	Until deleted by user or organisation; removed on tenant deletion	Including chunks and embeddings in Qdrant
Alert preferences and delivery logs	Duration of subscription or until deleted	
Correspondence	7 years	

10. Cookies

GRIP uses the following types of cookies:

- Strictly necessary cookies: required for platform operation, including authentication tokens and session state managed by Supabase Auth. These cannot be disabled without breaking the service.
- Analytics cookies: used to understand aggregate usage patterns. These are deployed only where you have given consent and may be withdrawn at any time via the cookie preferences panel.

We do not use third-party advertising cookies, cross-site tracking, or behavioural profiling. We do not sell or share your data with advertisers.

11. Data Security

We implement appropriate technical and organisational measures to protect personal data against unauthorised access, loss, destruction, or alteration:

- Encryption in transit (TLS 1.2+) across all platform components and API calls.
- Encryption at rest across Supabase, Qdrant, and S3 storage layers.
- Row-level security policies in PostgreSQL enforcing strict tenant isolation.
- Tenant and ownership filtering on all Qdrant vector search operations.
- Bearer token authentication and short-lived session tokens via Supabase Auth.
- Access controls limiting personnel access to personal data on a strict need-to-know basis.
- Data processing agreements in place with all third-party processors.

No internet transmission or electronic storage method is 100% secure. While we apply commercially reasonable measures, we cannot guarantee absolute security.

12. Children

GRIP is not intended for use by individuals under the age of 18. We do not knowingly collect personal data from children. If you believe we have inadvertently collected data from a minor, please contact us immediately and we will take prompt steps to delete it.

13. Changes to This Policy

We may update this Privacy Policy from time to time to reflect changes in our data practices, technology stack, or applicable law. We will notify registered users of material changes by email or by prominent notice within the platform. The version date at the foot of this policy indicates when it was last updated. Continued use of GRIP following notification of a material change constitutes acceptance of the updated policy.

14. Contact Us

For any questions about this Privacy Policy, to exercise your data rights, or to raise a concern, please contact:

GRIP — Guernsey Regulatory Intelligence Platform

Contact: via the GRIP website

Supervisory Authority: Office of the Data Protection Authority (ODPA), Guernsey — odpa.gg

---

Last updated: 19 March 2026