CaseFlow Automation Ltd

Information Security, Governance & Compliance Overview

Technical and organisational measures for client InfoSec due diligence

Version: 1.2 | Date: 25 June 2026 | Classification: External - Client Confidential

Contents

1. Executive Summary
2. About CaseFlow & Scope
3. Trust Model & Security Principles
4. Data Governance & Privacy (UK GDPR)
5. Security Architecture
6. Identity & Access Management
7. AI Governance & Model Risk Management
8. Monitoring, Logging & Incident Response
9. Business Continuity & Disaster Recovery
10. Third-Party Risk & Sub-Processors
11. Compliance & Control Mapping
12. Shared Responsibility
13. Appendices

1. Executive Summary

CaseFlow Automation Ltd provides a SaaS platform that helps UK credit hire professionals analyse insurer correspondence, draft responses, and receive strategic case guidance, all supported by AI cross-referenced against a curated legal knowledge base.

Security and privacy are built into the platform's architecture, not bolted on. Our approach is grounded in **`data minimisation`**, **`privacy by design`**, and **`honest framing`** of what our controls achieve and where their limits lie.

Key Facts

 Encryption: TLS in transit; AES-256 at rest; per-tenant data isolation via Row-Level Security.

 Access: Role-based access control, invite-only registration, concurrent session detection, email-verified authentication.

 Privacy: Three-layer PII protection — browser-side PDF processing, mandatory server-side masking gateway (now including bare/untitled name detection in field-label, salutation and sign-off contexts), and database-level masking applied at rest. No client data used for AI model training.

 AI Safety: Closed knowledge base, anti-hallucination controls, human-in-the-loop design. AI is decision-support only.

 Compliance: UK GDPR / DPA 2018 alignment; ICO registered (ZC013423); DPIA-ready architecture.

2. About CaseFlow & Scope

Service Description

CaseFlow Automation (CreditHire Assist) is a browser-based SaaS platform for UK credit hire professionals. Core features include:

- Correspondence analysis: identifying insurer arguments and cited authorities
- AI-assisted reply generation, cross-referenced against curated case law
- Live case advice with conditional strategic guidance
- Basic Hire Rate (BHR) challenge support with locality-based availability analysis
- Case law library, guidance notes, liability assessment, and response templates

Service Boundary

In Scope	Out of Scope
Web application (portal), server-side API functions, AI inference layer, data storage, PII masking gateway, authentication system	Client endpoints and devices, client on-premises networks, third-party integrations chosen by the client

Hosting Region

The platform is hosted in the **European Union** via managed cloud infrastructure. Data residency is EU-based.

3. Trust Model & Security Principles

- **Privacy by Design & Default:** Data minimisation is architectural. PDFs are processed locally, PII is masked before AI processing, and retention is purpose-limited.
- **Least Privilege:** Row-Level Security (RLS) enforces per-company data isolation at the database level. Users can only access their own company's data.
- **Defence in Depth:** Controls operate at multiple layers: browser, network, application, database, and AI prompt level.
- **Transparency:** We use honest language about what our controls achieve. We say "attempts to mask" rather than "guarantees removal". We document limitations alongside capabilities.

4. Data Governance & Privacy (UK GDPR)

4.1 Controller / Processor Roles

Role	Entity	Scope
Controller	Client organisation	Determines the purposes and means of processing claim data
Processor	CaseFlow Automation Ltd	Processes data on behalf of the client to provide the Service
Controller	CaseFlow Automation Ltd	Platform account data (user credentials, usage logs)

Lawful basis: Contractual necessity for providing the Service; legitimate interest in providing efficient legal support tools.

4.2 Data Classification

Class	Examples	Handling

Public	Marketing materials, published case law	Standard controls
Internal	Platform configuration, usage statistics	Access control, logging
Confidential	Correspondence content, case details, AI outputs	Encrypted at rest, RLS isolation, PII masking before AI processing
Restricted	Personal identifiers in uploaded correspondence	Masked before AI processing; original text stored in client's isolated partition only

4.3 Three-Layer Privacy Architecture

Layer 1: Local PDF Processing. Text-based PDF files are processed entirely in the user's browser using Mozilla's PDF.js (the same library used by Firefox). No PDF file is uploaded to any server. The user reviews extracted text before choosing to submit it for analysis. *Exception:* a scanned (image-only) PDF cannot be read by the browser; in that narrow case the file is uploaded to a private temporary storage bucket so it can be OCR'd, then deleted immediately after processing. A scheduled daily job removes any object older than 24 hours as a safety net.

Layer 2: PII Masking Gateway (in transit). Before any text is sent to the AI model, it passes through a mandatory server-side masking gateway that uses pattern-based rules tuned for UK claims data. Detected identifiers (titled names, **bare/untitled names in labelled or salutation/sign-off contexts**, emails, phone numbers, VRMs, NI numbers, driving licence numbers, IBAN, card numbers, sort codes, account numbers, policy/claim references such as POL, CLM, REF, CLAIM, POLICY, FILE, BHR, CHR, CHO, TPI and structured formats like 24-ABC-123456, street addresses) are replaced with neutral placeholders such as `[NAME_REDACTED]` or `[VRM_REDACTED]`. **UK postcodes are deliberately preserved** because they are material to basic-hire-rate and locality arguments; a separate output-side scrubber removes any postcode the AI introduces that was not present in the original input.

Layer 3: At-Rest Masking (in storage). The same masking patterns are reapplied via PostgreSQL `BEFORE INSERT/UPDATE` triggers using the database function `public.mask_pii()`. Records persisted to `correspondence`, `ai_replies`, `case_advice`, `preemptive_letters` and `taxi_risk_assessments` contain placeholders rather than raw identifiers, so saved history is masked even in the event of unauthorised database access.

The masking gateway is a mandatory processing step. There is no bypass path. Each invocation logs the count and categories of items masked. Original values are never logged.

⚠ **Honest Limitation:** Pattern-based detection is a strong first layer but not a perfect filter. Untitled names are detected only in high-precision contexts (field labels such as "Driver:" / "Renter Name:", salutations like "Dear …", and sign-offs); free-form prose containing only a bare name with no surrounding cue may not be detected. Novel reference formats and information that is only identifiable in combination (e.g. "the claimant's blue Ford Focus") are also out of scope. The masking gateway is an automated safety net, not a substitute for sensible data hygiene — users should send only the personal data the task actually needs. Anything that slips through is covered by the next layers: the AI provider operates under Zero Data Retention enterprise terms (data is not kept or used for training), and the data itself is deleted shortly after processing.

4.4 Data Retention & Deletion

CreditHire Assist is designed **not to retain personal case data**. Uploaded text and the outputs the tool generates are deleted immediately when the user is finished, on session close. A continuous scheduled purge runs as a backstop, so nothing is ever retained beyond 24 hours. Users save anything they need to keep to their own systems (a copy button is provided in-app) and re-run the task if required.

Data Type	Storage	Retention period
User account & profile	Encrypted, RLS-isolated	Life of account; deleted within 30 days of account closure
Uploaded text, AI replies, case advice, pre-emptive letters, taxi risk assessments and other case outputs	Encrypted at rest, PII-masked on write, RLS-isolated	Deleted immediately on session close; continuous scheduled purge ensures nothing is retained beyond 24 hours
Original (unmasked) PDF files	Not stored — processed in browser	N/A. Scanned-PDF OCR fallback files: deleted immediately after parsing; safety-net purge within 24 hours.

AI prompts (after masking)	Transient at the AI provider	Not retained for training; provider Zero Data Retention contract path
AI usage audit log (`ai_usage_log`)	Encrypted, masked metadata only (no PII payload)	12 months
Activity log (feature usage events)	Encrypted	12 months
Authentication / login attempt logs	Encrypted	12 months
Database backups (disaster recovery)	Encrypted, EU region	Uploaded files are never included in any backup. Standard disaster-recovery backups exist for the database, but because case content is deleted within 24 hours those backups hold essentially no personal case data, and once a record is deleted there is nothing to restore it from.
Policy acknowledgements	Encrypted	Life of account + 6 years (legal accountability)

We do not use client data to train AI models. Data is used solely to generate the specific output requested by the user.

Special category data: The Service is not intended for special-category data and users are asked not to submit it. We recognise that in credit hire work, health, injury or vulnerability information can incidentally appear in free-text correspondence. Where it does, the same safeguards apply (masking, data minimisation, encryption and immediate deletion), and as data controller the customer remains responsible for the Article 9 lawful basis.

4.5 Post-Upload Scrubbing of Residual PII

If, after upload, a user identifies that personal data slipped past the masking gateway, the affected record can be:

- Re-masked in place — the same record is updated, the at-rest trigger re-applies `mask_pii()`, and the previous unmasked text is overwritten in place (no historical row is retained for masked tables).
- Deleted entirely by the user from their history (correspondence, replies, advice, letters).
- Deleted from the AI provider — the AI provider operates under a zero-data-retention contract path: prompts and outputs are not stored beyond inference, so there is no model-side record to delete. We will provide written confirmation on request.

To request a bulk scrub or audit, contact info@caseflowautomation.co.uk; we aim to action within 5 working days.

4.6 Data Subject Rights

The platform supports the exercise of data subject rights including access, rectification, erasure, and portability. Users can delete their own correspondence and AI outputs. Company-level data deletion is supported on contract termination.

4.7 International Transfers

Platform data is hosted in the EU. Where AI model inference involves processing outside the UK/EEA, this occurs under the provider's enterprise data processing terms which include Standard Contractual Clauses (SCCs) and prohibit the use of input/output data for model training.

5. Security Architecture

5.1 Hosting & Infrastructure

- Hosted on managed cloud infrastructure with EU data residency
- TLS encryption for all data in transit
- AES-256 encryption for data at rest
- Managed database with automatic backups, point-in-time recovery, and encrypted backup storage
- Infrastructure managed by established cloud providers with SOC 2 Type II and ISO 27001 certifications

5.2 Application Security

- All AI interactions are mediated through server-side functions. Users never interact with the AI model directly.
- System prompts and knowledge base content are injected server-side; users cannot modify or view underlying instructions
- Input validation and sanitisation on all user-submitted data
- No API endpoint allows arbitrary prompts to the AI model
- Content Security Policy headers, HTTPS-only access

5.3 Network Security

- HTTPS (TLS) enforced for all connections
- API rate limiting on authentication and AI endpoints
- Server-side functions operate in isolated execution environments with no persistent state between invocations

6. Identity & Access Management

6.1 Authentication

Control	Implementation
Registration	**`Invite-only.`** Users must be invited by a platform administrator. No self-registration.
Email Verification	Required before first login
Password Security	Salted hashing (bcrypt); minimum complexity enforced
Session Management	Concurrent session detection. Only one active session per user. Automatic expiry for inactive sessions.
Password Reset	Secure token-based reset via email
Account Suspension	Administrators can immediately suspend user accounts

6.2 Role-Based Access Control (RBAC)

Roles are stored in a dedicated `user_roles` table, separate from user profiles, and enforced via database-level security functions.

Role	Access Level
Handler / User	Own company's data: analyse, draft, view history
Senior	Company-wide visibility of team activity
Manager	Company-wide visibility, usage statistics, team management
Platform Admin	User management, company administration, platform configuration

6.3 Data Isolation

Every database table containing client data enforces **Row-Level Security (RLS)** policies. These policies are evaluated at the database level on every query and cannot be bypassed by the application layer. A user from Company A cannot access, view, or modify data belonging to Company B, even if they manipulate API requests.

7. AI Governance & Model Risk Management

7.1 AI Use Cases & Human Oversight

AI is used exclusively for **decision-support**. It drafts, analyses, and suggests. It never makes legal decisions, sends correspondence, or takes autonomous action.

Feature	AI Role	Human Role
Correspondence Analysis	Identify insurer arguments & cited cases	Review, verify, decide response strategy
Reply Generation	Draft response cross-referenced against curated case law	Edit, approve, send
Case Advice	Provide conditional strategic guidance	Apply professional judgement
BHR Challenge	Analyse locality availability data	Review branch data and distances

Every AI output is presented as a **`draft requiring human review`**. Users must explicitly accept a disclaimer acknowledging this before accessing AI features.

7.2 Anti-Hallucination Controls

Control	Implementation
Closed Knowledge Base	AI can only cite cases from a curated, pre-loaded database. It is explicitly instructed not to cite anything outside this set.
Explicit System Prompts	Directives including *`"Do NOT invent case names"`*, *`"Only cite cases from the provided knowledge base"`*, *`"If no authority exists, say so."`*
Low Temperature	`temperature: 0.3` on all AI calls, reducing creative output and favouring factual responses
Knowledge Isolation	GTA claims receive only GTA protocol data; non-GTA claims receive case law only. No cross-contamination of authority sources.
Mandatory Limitation Language	When no relevant authority exists, the AI states this explicitly rather than speculating

> ⚠ **`No AI system can guarantee zero hallucination.`** These controls significantly reduce the risk, but users should always independently verify case law citations before relying on them.

7.3 Data Controls for AI

- **`PII masking before transmission:`** All text passes through the mandatory masking gateway before reaching the AI model
- **`No training on client data:`** Client correspondence and case details are never used to train or fine-tune AI models
- **`Prompt minimisation:`** System prompts include only the knowledge necessary for the specific task
- **`No persistent context:`** AI prompts are transient and are not stored or reused across sessions

- **`Server-side only:`** Users never interact with the AI model directly. All prompts are constructed and transmitted server-side.

7.4 Model Selection & Governance

- Models are selected centrally for their suitability for legal text analysis and are not configurable by end users
- AI model providers process data under enterprise data processing terms that prohibit training on input/output data
- Model changes are tested and validated before deployment

8. Monitoring, Logging & Incident Response

8.1 Logging & Audit Trail

Log Type	Content	PII Included?
Activity Log	Feature usage events with timestamps, user IDs, company IDs	No
PII Masking Log	Count and categories of items redacted per invocation	No. Original values are never logged.
Authentication Log	Login events, session creation, concurrent session detection	Email addresses (for identification)
Error Logs	Application and function errors	Designed to exclude PII

8.2 Monitoring

- Server-side function execution is monitored via managed cloud logging
- Usage dashboards available to company managers (team activity) and platform administrators (aggregate metrics)
- Automated health checks on AI inference endpoints

8.3 Incident Response

In the event of a suspected security or AI safety incident:

1. **Detect and Triage:** Identify the scope and severity of the incident using available logs and monitoring
2. **Contain:** Affected AI features can be disabled immediately at the platform level. User sessions can be terminated. API keys can be rotated.
3. **Investigate:** Audit logs allow identification of affected outputs, users, and timeframes
4. **Notify:** Affected clients are notified with details of the issue and recommended actions. Regulatory notification (ICO) within 72 hours where required under UK GDPR Article 33.
5. **Remediate:** Root cause analysis conducted; controls updated and tested before re-enabling affected features
6. **Review:** Post-incident review to update procedures and controls

8.4 AI-Specific Incident Actions

- Identify event type (fabricated case law, data leakage in output, unexpected model behaviour)
- Disable affected AI feature(s) at platform level
- Review prompt and masking logs to assess exposure
- Update system prompts, knowledge base, or masking rules as required
- Validate fix before re-enabling

9. Business Continuity & Disaster Recovery

Capability	Implementation
Database Backups	Automated daily backups with point-in-time recovery; encrypted backup storage
Data Redundancy	Managed by cloud infrastructure provider with multi-availability-zone resilience
Service Recovery	Stateless server-side functions can be redeployed rapidly; no single point of failure in the processing layer

Graceful Degradation	If AI inference is unavailable, the platform's non-AI features (case law library, templates, liability guide, guidance notes) remain operational
Data Portability	Client data can be exported on request

⚠ **Honest Limitation:** CaseFlow does not currently maintain a formalised Business Continuity Plan (BCP) or conduct scheduled DR exercises. Recovery capabilities are provided by the underlying managed cloud infrastructure. We are evaluating formal BCP documentation as the platform matures.

10. Third-Party Risk & Sub-Processors

We engage the following sub-processors to deliver CreditHire-Assist. Their processing is governed by data processing terms no less protective than those we owe our customers. The list is reviewed regularly and updated on material change.

Sub-processor	Service provided	Location
Supabase Inc. (via Lovable Cloud)	Managed PostgreSQL, authentication, file storage; hosted on AWS	EU (eu-west-2, Ireland)
Lovable GmbH	Application hosting and AI gateway layer	EU
Google LLC (Gemini, via Lovable AI Gateway)	AI inference under zero-retention enterprise terms	EU routing via Lovable; transfers under UK IDTA / SCCs
Cloudflare, Inc.	CDN, TLS termination, traffic proxying	Global edge, EU-preferred
Resend, Inc.	Transactional email	EU / US, under DPA

Lovable's published Data Processing Agreement, including the published sub-processor list, is available at trust.lovable.dev and is incorporated by reference. The hosting region for our project is EU (Europe, Ireland), evidenced at platform level. Transfers outside the EU at the AI inference layer are safeguarded through Standard Contractual Clauses at the

Lovable-to-Google contractual layer, with the formal IDTA / UK Addendum to the EU SCCs and Transfer Impact Assessment documentation available via Lovable's Trust Center on B2B compliance request.

11. Compliance & Control Mapping

11.1 Regulatory Alignment

Regulation / Framework	Status
UK GDPR / DPA 2018	Aligned: data minimisation, privacy by design, transparency, accountability, lawful basis documented
ICO Registration	Registered (ZC013423)
EU AI Act (Decision-Support)	Aligned: human-in-the-loop design, AI outputs labelled as drafts, disclaimers enforced
ISO/IEC 27001	Not certified. Controls are aligned with key domains (see mapping below). Formal certification is under consideration as the platform scales.
Cyber Essentials	Under consideration for future certification

11.2 Control Mapping (Against ISO 27001 Domains)

Domain	Controls Implemented	Evidence / Reference
Access Control	RBAC, RLS, invite-only registration, session management	Database policies; Admin panel
Cryptography	TLS in transit, AES-256 at rest, salted password hashing	Infrastructure configuration
Operations Security	Activity logging, PII masking logs, usage dashboards	Activity log table; Admin dashboards
Supplier Relationships	Limited sub-processor set; DPAs in place	Sub-processor register (§10)

Incident Management	Feature-level kill switches, audit trail, notification procedures	Incident response plan (§8.3)
Business Continuity	Automated backups, PITR, stateless architecture	Infrastructure provider capabilities
Data Protection	Three-layer PII protection (local processing, in-transit masking, at-rest masking), data minimisation, no training on client data	PII Masking Architecture document

12. Shared Responsibility

Area	CaseFlow Responsibility	Client Responsibility
Platform & Infrastructure	Secure hosting, encryption, monitoring, patching	-
Identity & Access	Authentication system, RBAC enforcement, session controls	Manage user invitations, remove leavers promptly, enforce strong passwords
Data Protection	Encryption, RLS isolation, PII masking, no-training policy	Classify data appropriately; redact/anonymise sensitive data before upload where possible
AI Usage	Anti-hallucination controls, closed knowledge base, disclaimers	Review all AI outputs before use; do not treat drafts as verified legal advice
Endpoints & Devices	-	Secure devices, keep browsers updated, use trusted networks
User Training	Platform guidance, in-app tooltips, onboarding materials	Ensure users understand the tool's purpose and limitations

13. Appendices

Appendix A: Related Documents

Document	Purpose
AI Safety & Security Policy	Detailed AI governance, anti-hallucination controls, prompt security, and user safeguards
How We Protect Your Data	Plain-English guide to the three-layer privacy architecture (local processing, in-transit masking gateway, and at-rest masking), with IT verification steps
Data & AI Summary	One-page summary of data processing and AI usage
Privacy Policy	Full privacy policy including lawful basis, rights, and contact details
PII Masking Architecture: Technical Summary	DPIA-ready technical reference for the PII masking gateway

Appendix B: Glossary

Term	Definition
RLS	Row-Level Security. Database-level access control that restricts which rows a user can read or modify.
PII	Personally Identifiable Information. Data that can identify a natural person.
RBAC	Role-Based Access Control. Access permissions determined by assigned roles.
GTA	General Terms of Agreement. ABI protocol governing credit hire rates.
BHR	Basic Hire Rate. Insurer argument challenging hire charges based on local availability.
PITR	Point-in-Time Recovery. Ability to restore a database to a specific moment.
SCCs	Standard Contractual Clauses. EU-approved mechanism for international data transfers.
DPA	Data Processing Agreement. Contract governing how a processor handles personal data.

DPIA	Data Protection Impact Assessment. Risk assessment required for high-risk processing.

Appendix C: Contact and Enquiries

For security-related enquiries, DPIA collaboration, or to request further technical detail:

Email: info@caseflowautomation.co.uk

ICO Registration: ZC013423

Website: caseflowautomation.co.uk

This document is designed to be **honest and specific** about our security posture. We describe what we have implemented, acknowledge where we have limitations, and commit to continuous improvement as the platform matures. **Trust is earned through transparency, not theatre.**
