

# Trust & Security Pack

Security, data governance and AI governance in one document

Version: 4.0 | Date: 6 October 2026 | Classification: External - Client Information

---

## 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
14. AI Safety & Security Policy
15. Technical Overview

One consolidated reference for operator IT, security, compliance and procurement teams. It brings the security overview, AI policy, technical overview and data summary together without repeating four separate document covers.

This pack describes TotalSettle, not the different AI tools or retention models used by other CaseFlow products. The signed DPA and agreed service terms remain the contractual authority.

## 1. Executive Summary

TotalSettle is a digital total loss settlement portal for UK motor claims operators. Handlers review engineer report figures, prepare an offer and issue a branded settlement link. Claimants verify their details before viewing, accepting or disputing the offer.

### Key facts

Encrypted storage and transport. Server-side organisation isolation. Registration and postcode verification before case disclosure. Personal identifiers redacted in the engineer-report text path. Handler review before issue. Personal data removed on completion, with a 30-day backstop for open cases.

## 2. About CaseFlow & Scope

CaseFlow Automation Ltd publishes TotalSettle. Registered office: 7-9 Macon Court, Crewe, CW1 6EA. ICO registration: ZC013423.

In scope: the operator portal, claimant settlement portal, report extraction, evidence storage, email and SMS delivery, server-side processing, database access controls and retention routines. Client devices, local networks and client-selected integrations remain outside our operational boundary.

## 3. Trust Model & Security Principles

**Privacy by design:** collect the fields needed for settlement and remove personal data when it is no longer needed. **Least privilege:** authorise access server-side. **Defence in depth:** combine authentication, tenant isolation, verified claimant grants and validation. **Transparency:** describe limitations as well as safeguards; redaction is not a guarantee that every identifier is recognised.

## 4. Data Governance & Privacy (UK GDPR)

**Roles.** The operator is Data Controller for claim data. CaseFlow Automation Ltd is Data Processor under the operator's instructions and a signed DPA. Operators remain responsible for their lawful basis, privacy notices and handling data subject requests.

**Data classes.** Published service information is public. Configuration and usage metadata are internal. Claim figures and evidence are confidential. Claimant identifiers and contact details require restricted, authorised access.

**Data flow.** Browser PDF extraction is used where possible. Engineer report text passes through server-side identifier redaction before AI extraction. Names, vehicle registration, postcode, email, phone, National Insurance number, UK driving licence, IBAN and card identifiers are replaced with opaque placeholders where detected. Original values are restored locally into the case record. Unlike some sister products, TotalSettle retains the identifiers needed to run a live settlement; it does not claim that all live case fields are masked at rest.

**Important limitation.** Pattern-based redaction may miss unexpected formats or identifiers in free-form text. Where a scanned or image-only report cannot be read locally, the original PDF is sent to the AI provider for extraction. This path cannot be redacted in-place and its use is logged per case.

**Live-case retention.** Only the extracted fields needed for the settlement are held. Raw engineer PDFs are not retained beyond extraction. Claimant evidence is held while the case remains open.

**Completion and expiry.** On acceptance, claimant identifiers, engineer comments, finance company and the settlement token are removed; evidence rows and files are deleted; related email and SMS log personal fields are cleared. Open cases are anonymised at 30 days from creation and marked expired, with an operator warning 7 days beforehand. An anonymised statistical record remains for reporting. Every purge is audited. Tenant offboarding removes the organisation's data, subject to legal hold.

**Residency and rights.** Database storage is in the EEA, AWS eu-central-1 (Frankfurt). UK-EEA transfers rely on applicable adequacy arrangements. Operators can request assistance with access, rectification, erasure and portability through the agreed DPA process.

## 5. Security Architecture

The browser application is served from a CDN. Managed Postgres stores case records with row-level organisation isolation. Server-side functions handle report extraction, authorised outbound communication and lifecycle operations. Object storage holds evidence and organisation logos. Data is encrypted in transit and at rest.

Private provider credentials stay server-side. Authorisation precedes case-specific send checks. Case creation does not require a postcode, but email and SMS settlement sends are refused when no postcode is on file. Input validation, HTML escaping and constrained external requests reduce injection and request-abuse risks.

## 6. Identity & Access Management

Operator users authenticate with managed accounts. Password reset and team invitations use dedicated flows. Role assignments are held separately from profiles, with server-side role checks and organisation membership controls. Authentication capabilities and assurance requirements should be confirmed during procurement rather than inferred from this summary.

Claimants do not need an account. A high-entropy link token plus matching vehicle registration and postcode is required before any case details are returned. Failed verification uses generic responses and throttled retries. Successful verification issues a short-lived access grant. Subsequent evidence and response operations use the grant rather than an unverified token.

An expired grant sends the claimant back to verification. Grants are revoked on successful acceptance and during personal-data purge. Evidence additions are limited to open cases. Settlement links have no independent time-based expiry; access ends when the case closes or is purged and the data and token are removed.

## 7. AI Governance & Model Risk Management

Google Gemini is accessed through the Lovable AI gateway for structured engineer-report extraction and a short customer summary. The provider use is subject to no-training terms. This pack does not independently certify provider-side retention periods; the current contractual position should be confirmed in procurement.

AI supports a trained handler, never makes the final settlement decision. Extracted figures and summaries require review before send. Structured fields and validation constrain outputs, but accuracy cannot be guaranteed. A release gate blocks new AI communication unless it is a written, approved exception. Existing exceptions cover report reading and supporting extraction, plus an admin-only marketing voiceover tool that does not process case data.

## 8. Monitoring, Logging & Incident Response

Handler and claimant actions are recorded with timestamps and identifying audit context. Redaction logs record counts and categories, not the original identifier values. Scanned-report fallbacks are recorded per case. Email and SMS sends have delivery records; their personal fields are cleared by the case purge. Purge audit records and automated health checks provide evidence that retention routines ran.

Incident handling follows detection and triage, containment, investigation, notification, remediation and review. Affected controllers are informed through the agreed incident process. Controllers assess regulatory notification duties, including the UK GDPR 72-hour deadline where applicable. For a suspected incident, contact [info@caseflowautomation.co.uk](mailto:info@caseflowautomation.co.uk).

## 9. Business Continuity & Disaster Recovery

Managed database backups and point-in-time recovery support disaster recovery. Deletion from the live service does not mean immediate removal from every backup: purged data falls out as backups rotate through their recovery window. Backup windows, recovery objectives and restoration procedures should be confirmed in the service agreement. Disaster recovery procedures are documented and reviewed.

## 10. Third-Party Risk & Sub-Processors

The managed cloud platform provides database, identity, storage and server-side execution. Resend supports transactional email. FireText Communications Ltd supports SMS under a data processing agreement with UK-only message routing. SMS carries the settlement link and message context, not valuation or claim details. Google Gemini is accessed through the Lovable AI gateway for approved extraction and summary tasks.

The current contractual sub-processor list, processing locations and transfer safeguards are available on request. Do not assume all third-party inference or delivery processing shares the database's Frankfurt location.

## 11. Compliance & Control Mapping

**UK GDPR / Data Protection Act 2018:** purpose-limited processing, minimisation, controlled access, deletion and controller assistance under the DPA. **Access control:** authenticated operator membership and verified claimant grants. **Accountability:** action records and purge audit evidence. **AI governance:** redaction, documented exceptions and human review.

The platform supports FCA Consumer Duty evidencing through clear offers, review, dispute handling and an audit trail. This is not FCA or Ombudsman endorsement. ICO registration is not certification of every control. No sister-site certification, testing result or assurance claim is transferred to TotalSettle without separate evidence.

## 12. Shared Responsibility

**CaseFlow:** operate platform controls, enforce tenant boundaries, manage approved processing and execute retention routines. **Operator:** maintain lawful instructions, invite appropriate users, protect accounts and devices, check extracted figures, keep a valid postcode before sending, review settlement fairness and retain any records it must hold in its own system.

Claimants verify registration and postcode, review the offer and choose whether to accept or dispute. Operators handle disputes; AI cannot accept an offer or override a claimant decision.

## 13. Appendices

**Supporting documents:** AI Data Handling, Data & AI: Plain English Guide and How We Protect Your Data. The privacy policy and Terms of Use remain separate legal documents. Legacy document links are retained for compatibility, but this consolidated pack is the current due-diligence reference.

**Available on request:** signed DPA, current sub-processor list, DPIA and security questionnaire responses. Contact [info@caseflowautomation.co.uk](mailto:info@caseflowautomation.co.uk). Public AI reference: <https://totalsettle.com/ai-data-handling>.

This document is provided for information and does not constitute legal advice or replace signed contractual terms.

## 14. AI Safety & Security Policy

Approved AI tasks are engineer-report field extraction and a short plain-English portal summary. Prompts are maintained server-side. Report content is untrusted input, not an instruction source. The intended response is a structured set of fields; numeric and field validation and handler review are used before issue.

The text route must apply the shared redaction and local restoration process. Placeholders preserve the association between extracted identifiers and the locally held originals. The image-only PDF fallback is an explicit, logged exception, not a promise of universal redaction. Never describe it as fully masked.

Model outputs can be inaccurate, incomplete or influenced by malicious document content. Handlers must check source figures and customer wording. A claimant may dispute and provide evidence. Changes to approved AI processing require review, an approved release-gate exception where applicable, and an update to the published processing description.

AI processing is not used for final settlement decisions. No-training terms apply to approved model processing. Provider retention and transfer terms should be checked against the current DPA and sub-processor schedule.

## 15. Technical Overview

Frontend: React and TypeScript with Vite and Tailwind. Database: managed Postgres with row-level security. Server-side processing: managed Deno functions. PDF extraction: browser PDF.js where possible, with shared server-side redaction, structured Gemini extraction and local reinsertion.

Claimant verification returns safe fields and a short-lived access grant. Only a hash of the grant is stored server-side. Verified response and evidence routines resolve the case through the grant. Acceptance and purge revoke associated grants. Legacy unverified case/evidence accessors are not used.

Retention combines event-driven purge with scheduled completed-case and abandoned-case safety nets, operator warnings, durable purge audit tracking and automated checks. Offboarding uses an audited tenant-deletion routine. Evidence, email and SMS personal fields participate in the same lifecycle.

Operational implementation names are not a substitute for assurance evidence. Request the current DPA, sub-processor list and supporting security information for formal due diligence.