

AI Safety & Security Policy

How AI is used inside TotalSettle, with redaction, oversight and logging.

Version v3.2 | Last reviewed: 24 July 2026 | Owner: CaseFlow Automation Ltd

Scope

This policy governs how AI is used inside TotalSettle, including data handling, anti-hallucination controls, prompt security, and human oversight. It applies to all AI calls made by the platform.

Approved AI provider

TotalSettle currently uses Google Gemini, accessed through the Lovable AI gateway. The provider does not retain data submitted through the gateway and does not use it to train any model. Processing is performed in supported regions only. Changes to the approved provider are subject to a security and DPIA review and are reflected on the public AI data handling page before going live.

What is sent to AI

- Text extracted from the engineer report, after server-side PII redaction.
- Structured field prompts asking the model to return valuation, salvage, category, vehicle registration, claimant identifiers and reference.
- A request to draft a short plain-English summary for the customer portal.

PII redaction

Before the engineer report text leaves our system, identifiers are replaced with opaque placeholders such as NAME_1 or VRM_1. Categories redacted include names (titled and labelled), vehicle registration, postcode, email, phone, National Insurance number, UK driving licence number, and IBAN or card numbers. The AI provider sees only the placeholders in the text path. The original values are restored locally on the way back, into the structured case record. Per-case redaction counts (by type, never values) are logged so the control can be evidenced.

Image-only PDF fallback

Some engineer reports are scanned or image-only and defeat in-browser text extraction. In that case the original PDF is sent to the AI provider so the fields can still be read. This fallback cannot be redacted in-place. Every use of the fallback is logged per case so frequency can be reviewed and the practice can be tightened if required.

Anti-hallucination controls

- Structured output schema. The model is required to return named fields, not free prose.
- Numeric values are validated against expected ranges and field types before being saved.
- Confidence indicators surface to the handler when extraction is uncertain.

- Handlers review extracted figures and the generated customer summary before send.
- The customer-facing portal shows the source figures and a documented dispute route.

Prompt security

- System prompts are version-controlled and reviewed before deployment.
- Engineer report text is treated as untrusted input. Instructions inside the document are ignored by the model under explicit system instruction.
- Placeholders such as NAME_1 are required to be preserved verbatim. Outputs that lose the placeholders are rejected and reprocessed.

Human oversight

AI is decision support, not the decision maker. A trained handler reviews extracted fields and the customer summary before the settlement link is issued. The customer can accept or dispute. Disputes go back to the handler with full audit history.

Logging and review

- Each AI extraction is logged against the case with timestamp and model identifier.
- Per-case redaction summary (counts only) is recorded.
- Use of the image-only fallback is recorded and reviewed.
- Logs are retained for the life of the case. When a case's personal data is purged (on completion, or at 30 days for abandoned cases), email logs are redacted and the purge itself is recorded in a dedicated audit trail.

CaseFlow Automation Ltd is the holding company that publishes TotalSettle. Registered in England and Wales. ICO registration: ZC013423. Registered office: 7-9 Macon Court, Crewe, CW1 6EA. This document forms part of the TotalSettle policy pack and is provided for information. It does not constitute legal advice. For the signed DPA, sub-processor list, or DPIA, contact info@caseflowautomation.co.uk.