

# PII Masking Architecture

Technical Summary for DPIA / Compliance Reference

|                  |                                                |
|------------------|------------------------------------------------|
| Platform         | CaseFlow Automation (caseflowautomation.co.uk) |
| Last Updated     | 30 June 2026                                   |
| Classification   | Internal - Compliance                          |
| Document Purpose | DPIA / Compliance Reference                    |

## 1. Overview

CaseFlow Automation processes insurer correspondence and case details using AI to generate legal analysis and draft responses. To support UK GDPR data-minimisation principles (Article 5(1)(c)), the platform implements a **three-layer privacy architecture**:

| Layer                                      | Location                    | Purpose                                                                                                                                                                                                                                                                                                                                                                                                                |
|--------------------------------------------|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Layer 1 - Local Processing                 | User's browser              | Text-based PDFs are extracted entirely client-side; the file never reaches our servers. A server-side fallback is used only when a PDF is a scanned image (no embedded text), in which case the file is uploaded to a private temporary bucket for OCR and removed immediately after processing.                                                                                                                       |
| Layer 2 - PII Masking Gateway (in transit) | Server-side (Edge Function) | Regex-based pattern matching detects and replaces personal identifiers with neutral placeholders before text is transmitted to the AI model.                                                                                                                                                                                                                                                                           |
| Layer 3 - At-Rest Masking                  | Database (triggers)         | The same masking patterns are applied via PostgreSQL triggers ( <code>public.mask_pii()</code> ) on insert/update to <code>correspondence</code> , <code>ai_replies</code> , <code>case_advice</code> , <code>preemptive_letters</code> , and <code>taxi_risk_assessments</code> . Saved history therefore contains placeholders, not raw identifiers, mitigating impact in the event of unauthorised database access. |

### Temporary File Retention

The `temp-pdfs` storage bucket is private. Files are deleted by the application immediately after OCR completes. A scheduled `pg_cron` job ( `purge-temp-pdfs-daily` ,

03:00 UTC) acts as a safety net, removing any object older than 24 hours. In typical operation the bucket is empty.

## 2. Layer 1 - Local PDF Processing

- PDF files are processed using `pdfjs-dist` (Mozilla's PDF.js) directly in the user's browser.
- Text is extracted locally and displayed in the UI for the user to review.
- **No PDF file or extracted content is transmitted to any server during this step.**
- The user then chooses to submit the text for AI analysis, at which point Layer 2 applies.

## 3. Layer 2 - PII Masking Gateway

### 3.1 Architecture

The masking gateway ( `piiRedactor.ts` ) is a shared server-side module invoked by every AI-facing Edge Function before any text is sent to the language model. It operates as a mandatory processing step on the server.

### 3.2 Data Categories Detected

The gateway attempts to detect and mask the following UK-specific identifier types:

| #  | Category                   | Placeholder                   | Example Patterns                                                                                                                                                                                                     |
|----|----------------------------|-------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1  | Email addresses            | [EMAIL_REDACTED]              | Standard email format                                                                                                                                                                                                |
| 2  | National Insurance Numbers | [NINO_REDACTED]               | AB 12 34 56 C                                                                                                                                                                                                        |
| 3  | UK Driving Licence Numbers | [DRIVING_LICENCE_REDACTED]    | 16-character DVLA format                                                                                                                                                                                             |
| 4  | Vehicle Registration Marks | [VRM_REDACTED]                | Current, prefix, suffix, dateless - requires $\geq 2$ digits to avoid car model false positives                                                                                                                      |
| 5  | UK Postcodes               | <i>Not masked (preserved)</i> | Deliberately preserved because postcodes 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.  |
| 6  | UK Phone Numbers           | [PHONE_REDACTED]              | Landline, mobile, +44 international                                                                                                                                                                                  |
| 7  | Dates of Birth             | [DOB_REDACTED]                | Contextual only (preceded by DOB/born)                                                                                                                                                                               |
| 8  | Bank Account Numbers       | [ACCOUNT_REDACTED]            | Account number with contextual prefix                                                                                                                                                                                |
| 9  | Policy / Claim References  | [REF_REDACTED]                | Markers include POL, CLM, REF, CLAIM, POLICY, FILE, BHR, CHR, CHO, TPI, plus structured formats such as 24-ABC-123456 and ABC/2025/12345                                                                             |
| 10 | Street Addresses           | [ADDRESS_REDACTED]            | House number + street name (conservative)                                                                                                                                                                            |
| 11 | Person Names               | [NAME_REDACTED]               | Titled names (Mr/Mrs/Ms/Miss/Dr/Prof + surname) are always masked. Bare/untitled names are masked only in high-precision contexts: field labels (e.g. "Claimant:", "Driver:", "Insured:"), salutations ("Dear ..."), |

|  |  |  |                                                                                                                                                                                       |
|--|--|--|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  |  |  | sign-offs, and "Re <name>" in email subjects. Guards protect case citations and company names. A bare first name and surname in free prose with none of those cues may not be caught. |
|--|--|--|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

### 3.3 Processing Order

Patterns are applied in a specific sequence (most specific first) to prevent partial matches or false positives:

1. Emails → 2. NINOs → 3. Driving Licences → 4. VRMs → 5. Phone Numbers → 6. DOBs → 7. Bank Details → 8. References → 9. Addresses → 10. Titled Names → 11. Bare Personal Names. UK postcodes are deliberately not part of the masking sequence; they are preserved for locality/BHR analysis, and a separate output-side scrubber removes any postcode the AI introduces that was not present in the original input.

### 3.4 Logging

Each invocation logs the count and types of items masked (e.g., PII redaction: 3 items redacted (types: EMAIL, VRM, REF) ). No original PII values are logged.

## 4. Integration Points

The PII Masking Gateway is invoked in the following backend functions:

| Function                 | Purpose                         |
|--------------------------|---------------------------------|
| analyze-correspondence   | Insurer correspondence analysis |
| generate-reply           | AI-generated reply drafting     |
| case-advice              | Live case strategic guidance    |
| generate-argument-letter | Argument letter generation      |

**All** text submitted to the AI model passes through the gateway. There is no bypass path.

## 5. GTA vs Case Law Data Separation

As of 13 February 2026, the platform enforces strict separation between GTA protocol data and case law:

| Claim Type                      | Knowledge Provided to AI                                | Case Law Included?                          |
|---------------------------------|---------------------------------------------------------|---------------------------------------------|
| <b>GTA Claim</b> (ticked)       | GTA sections, hire rate tables, EV-to-ICE mappings only | <b>No</b> - case law is explicitly excluded |
| <b>Non-GTA Claim</b> (unticked) | Case law database                                       | <b>Yes</b> - full case law knowledge base   |

This ensures GTA claims receive advice grounded solely in GTA protocol and rate tables, while non-GTA claims receive case law citations. The separation prevents cross-contamination of authority sources and reduces hallucination risk.

### EV Handling Under GTA

Electric vehicles not listed in standard GTA rate tables are mapped to ICE equivalent groups with mandatory explanatory notes. Accident dates are preserved as non-identifying data.

## 6. Anti-Hallucination Controls

| Control                                 | Implementation                                                                                                                                                                                                                                                                                  |
|-----------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Low Temperature</b>                  | temperature: 0.3 on all AI calls                                                                                                                                                                                                                                                                |
| <b>Explicit System Prompts</b>          | "Do NOT invent case names", "Only cite cases from the provided knowledge base"                                                                                                                                                                                                                  |
| <b>Knowledge Isolation</b>              | GTA claims receive no case law; non-GTA claims receive no GTA rate data                                                                                                                                                                                                                         |
| <b>Final Output Scrubbing (Layer 3)</b> | Active — <code>citationScrubber.ts</code> runs on all AI outputs. Strips fabricated or unverified case citations while preserving legitimate legal references. Applied across all six edge functions (Analysis, Reply, Case Advice, Liability Assessment, Argument Letter, Pre-Emptive Letter). |

## 7. Limitations & Honest Framing

The platform UI and Privacy Policy use deliberately honest language:

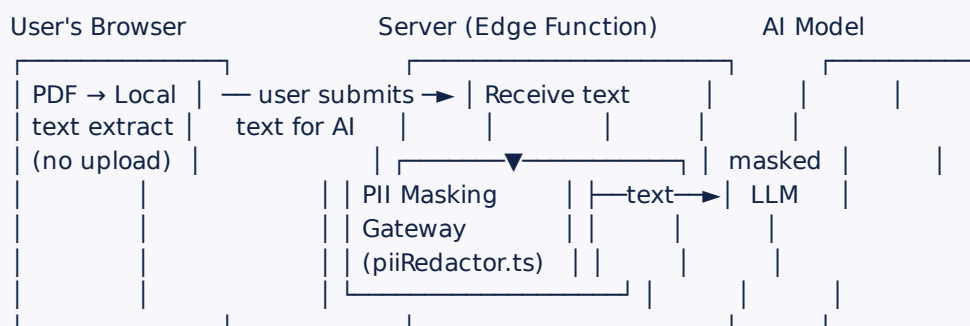
- **"attempts to mask"** - acknowledges that pattern-based detection is a strong first layer rather than a perfect filter, so users are asked to send only the personal data the task actually needs.
- **"where detected"** - frames the capability as best-effort pattern matching, not absolute removal.

- **Automated safety net** - the masking gateway operates as a mandatory server-side step, removing the requirement for users to manually redact documents before upload.

### Known Limitations

- **Names:** Titled names (Mr, Mrs, Dr, etc.) are always masked. Bare/untitled names are masked only in high-precision contexts: field labels (e.g. "Claimant:", "Driver:", "Insured:"), salutations ("Dear ..."), sign-offs, and "Re <name>" in email subjects. Guards protect case citations and company names. A bare first name and surname in free prose with none of those cues may not be caught. Coverage of unusual or ambiguous names cannot be guaranteed.
- **Novel Formats:** Unusual or non-standard reference numbers may not match existing patterns.
- **Contextual Data:** Information that is only personally identifiable in combination (e.g., "the claimant's blue Ford Focus") is not masked, as it requires semantic understanding beyond regex capability.

## 8. Data Flow Summary



## 9. Regulatory Alignment

| Principle                               | Implementation                                                                          |
|-----------------------------------------|-----------------------------------------------------------------------------------------|
| <b>Data Minimisation</b> (Art. 5(1)(c)) | Three-layer approach: local processing, in-transit masking gateway, and at-rest masking |
| <b>Privacy by Design</b> (Art. 25)      | Masking gateway is mandatory in the processing pipeline; no bypass                      |
| <b>Transparency</b> (Art. 13/14)        | UI wording and Privacy Policy describe the approach honestly                            |
| <b>Accountability</b> (Art. 5(2))       | Source code is auditable; server logs record masking activity without storing PII       |

## 10. Recommendations for Further Hardening

- **Periodic Pattern Review:** Review and update regex patterns quarterly as new identifier formats emerge.
- **User Training:** Reinforce during onboarding that the masking layer is a safety net, not a substitute for careful data handling.
- **DPIA Integration:** Reference this document in the platform's Data Protection Impact Assessment.

---

*This document describes the technical architecture as implemented by CaseFlow Automation. It does not constitute legal advice. Organisations should obtain independent legal counsel for DPIA completion and regulatory compliance assessments.*