

How It Works - A Plain English Guide

Everything you need to know, without the jargon

Last Updated: 25 June 2026

We know that "AI", "data privacy", and "GDPR" can sound intimidating. This guide explains what CaseFlow actually does, how it keeps your information safe, and what you need to know - all in plain English.

What Does CaseFlow Do?

- CaseFlow helps credit hire teams deal with insurer correspondence faster and with greater confidence. You paste in (or upload) a letter from an insurer, and the system analyses the legal arguments they're making, identifies which court cases they're relying on, and helps you draft a professional response.

Think of it as a very well-read assistant that's studied all the key credit hire case law - but one that always needs you to check its work before anything goes out the door.

What AI Model Does It Use?

- CaseFlow uses **Google's Gemini 3 Pro** - one of the most capable AI models available - for tasks that need careful legal reasoning, like case advice and drafting argument letters. For quicker tasks like initial analysis, it uses **Gemini 3 Flash**, which is faster but still very capable.

These are the same models used by major businesses worldwide. We don't use an experimental or unproven system.

Does the AI Make Things Up?

 This is the biggest risk with AI in a legal context - it's called "hallucination", where the AI invents a case name or citation that doesn't exist. We take this extremely seriously and have built a **three-layer protection system** that every AI output must pass through before it reaches you:

- **Layer 1 - Prompt rules:** The AI is given strict instructions that it is *"strictly prohibited from citing, referencing, or mentioning any case law that does not appear in the knowledge base."* If it can't find authority, it must say so - not guess.
- **Layer 2 - Citation validator:** After the AI generates its response, every case name it cites is automatically cross-referenced against our verified case law database. Any citation not found in our database is stripped from the output before you see it.
- **Layer 3 - Text-body scrubber:** A separate scanner reads through all the prose text (summaries, reasoning, strategic advice) looking for any remaining "X v Y" case-name patterns. If it finds one that isn't in our verified database, it replaces it with a warning placeholder.

Think of it like a three-checkpoint security gate. First, the AI is told what it's allowed to cite. Then, its citations list is checked against our records. Finally, the entire text is scanned for any case names that slipped through. All three checks happen automatically before anything reaches your screen.

That said, no AI is perfect. We always recommend checking case citations independently before using them. The system is a starting point, not a final answer.

What Happens to My Documents?

 **Your PDFs never leave your computer** - with one narrow exception. When you upload a normal text-based PDF, it's read and converted to text right there in your web browser. The actual file is never sent to our servers.

The exception is a **scanned (image-only) PDF**: because the browser can't read the text, the file is briefly uploaded to a private temporary store so it can be OCR'd, then deleted immediately. As a safety net, anything left over is automatically wiped within 24 hours. In day-to-day use this should virtually never happen.

Only the text content is sent for analysis - and even then, it goes through our privacy filter first (more on that below).

It's like photocopying a letter at your desk and only sending the copy - the original stays with you, and we blank out personal details on the copy before passing it along.

How Can You Be Sure?

We understand that "trust us" isn't always enough - particularly if your IT team or a client wants proof. The short answer: any IT professional can watch your browser's network traffic while a PDF is being loaded. They will see **zero outbound requests** at that stage. The only call to our servers happens *after* you click Analyse, and it carries extracted text - not the original file. The document genuinely never moves.

This isn't just our policy - it's how the underlying technology works. PDF reading runs entirely inside the browser tab, using the same open-source library that Firefox uses to display PDFs natively. It's physically incapable of uploading the file.

How Is Personal Information Protected?



Before any text is sent to the AI, it passes through our **PII masking gateway**. "PII" stands for Personally Identifiable Information - things like names, email addresses, phone numbers and registration plates.

Our system attempts to detect these and replaces them with placeholders. So instead of seeing "Mrs Lisa Chen", the AI sees "[NAME_REDACTED]". Instead of "john.smith@example.com", the AI sees "[EMAIL_REDACTED]". Instead of "AB12 CDE", it sees "[VRM_REDACTED]".

Here's what we try to catch:

-  Names with a title (Mr, Mrs, Ms, Miss, Dr, Prof + surname) are always masked. Names without a title are masked when they appear in obvious places like field labels ("Claimant:", "Driver:", "Insured:"), salutations ("Dear ..."), sign-offs, and "Re <name>" in email subjects. A bare first name and surname buried in free prose with none of those cues may not be picked up.
-  Email addresses
-  Phone numbers
-  Vehicle registration plates
-  National Insurance numbers
-  Driving licence numbers
-  Bank account details
-  Policy and claim reference numbers - markers like POL, CLM, REF, CLAIM, POLICY, FILE, BHR, CHR, CHO, TPI, plus structured formats such as 24-ABC-123456 and ABC/2025/12345
-  Street addresses

Postcodes are deliberately kept. UK postcodes matter to locality and basic-hire-rate arguments, so we leave them in the text the AI sees. As a safeguard, an output check then removes any postcode the AI tries to introduce that was not in your original text.

An honest note: Automated pattern-based masking is a strong first layer, but it isn't a perfect filter. That's why we also ask you to send only the personal data the task actually needs. Anything that slips through is covered by the next layers: our AI provider operates under Zero Data Retention terms (your data is not kept or used to train models), and the data itself is deleted shortly afterwards on session close, with a continuous purge as a backstop so nothing is retained beyond 24 hours.

Masked Twice - In Transit and At Rest

The masking runs in **two places**: once before text is sent to the AI, and again at the database level when results are saved. That means the history we keep on our servers (your past analyses, replies, advice and letters) already contains

placeholders like [NAME_REDACTED] rather than real names or contact details. If our database were ever accessed without authorisation, the personal information has already been stripped out.

Is My Data Used to Train the AI?

- ☐ **No. Absolutely not.** We do not use your data to train AI models. Your correspondence and case details are used solely to generate the specific response you asked for - nothing more. The AI provider's terms also prohibit them from using your data for training.

Can Other Companies See My Data?

- ☐ **No.** Every company on CaseFlow operates in its own completely separate space. Your correspondence, case advice, templates, and history are invisible to other companies. This is enforced at the database level - it's not just a setting that could be accidentally turned off.

Think of it like separate filing cabinets in separate locked rooms. Each company has their own room, and there's no shared key.

What About GDPR?

-  **GDPR** (the UK's data protection law) requires organisations to only collect and process the minimum personal data necessary, to protect it properly, and to be transparent about what they do with it. Here's how we line up:

- ☐ **Data minimisation:** PDFs stay on your device. Text is masked before reaching the AI.

- ☐ **Privacy by design:** The masking gateway is mandatory - there's no way to bypass it.

- **Transparency:** We tell you exactly what we do (you're reading it now).
- **Your rights:** You can request access to, correction of, or deletion of your data at any time.
- **ICO registered:** We're registered with the Information Commissioner's Office (ZC013423).

Is It Legal Advice?



No. CaseFlow is a decision-support tool. It helps you draft responses and understand legal arguments, but it's not a substitute for qualified legal advice. Every output should be reviewed by someone who understands the specific case before it's used.

We make this clear throughout the platform - you'll see reminders when you first log in, on every page that uses AI, and alongside every generated response.

The Quick Summary

- **Your documents stay on your device** - we only see the text, never the file (scanned PDFs are the one narrow exception, and are deleted within 24 hours).
- ♥ **Personal details are filtered out** before the AI sees anything, and again before anything is saved to our database.
- **We use Google's Gemini 3 Pro** - a leading, enterprise-grade AI model.
- **Your data is never used to train the AI** - full stop.
- **Your data is completely separate** from every other company on the platform.
- ⚙️ **3-layer protection against made-up case law** - prompt rules, citation validator, and text-body scrubber all run automatically.

⚖️ **We're GDPR compliant** and registered with the ICO.

⚠️ **It's a tool, not a lawyer** - every output needs human review.

How We Handle Your Data

Behind the scenes, your PDFs stay in your browser - they're never uploaded to our servers. Before any text reaches the AI, we mask common personal details (emails, phone numbers, registration plates, and more) so the model works on cleaned-up text, not raw data. Postcodes are deliberately kept because they matter to locality arguments, and an output check stops the AI inventing a postcode that was not in your original text. And we never use your data to train AI models - full stop.

For the complete picture, see our [How We Protect Your Data](#) page.

Got questions? If anything in this guide isn't clear, or you'd like more detail on any point, get in touch at info@caseflowautomation.co.uk.

For the full technical details, see our [How We Protect Your Data](#) and [AI Safety & Security Policy](#) documents.

CaseFlow Automation Ltd - Legal Processes. Streamlined.

This guide is intended to help non-technical users understand how the platform works. It does not constitute legal advice.

See also: [How We Protect Your Data](#) | [AI Safety & Security Policy](#) | [Privacy Policy](#) | [Data & AI Summary](#)