



5 questions your AI-chat controls should answer

A two-page review checklist for UK law firms

An AI policy is only the starting point. A practical control should define where AI chat is used, act before confidential material crosses the boundary, and leave reviewable evidence without creating a second store of employee prompts.

1

Which AI-chat services and workflows are actually in scope?

WHY IT MATTERS

The Law Society recommends defining the purpose, use cases, input data and oversight for generative AI. A policy with no named owner, approved services or permitted data classes is difficult to apply consistently.

EVIDENCE TO LOOK FOR

- An approved-services register covering consumer accounts, enterprise copilots and browser-based chat.
- Permitted use cases, prohibited data classes, an accountable owner and a documented review cadence.
- A clear route for exceptions, incidents and questions from fee-earners.

2

What stops confidential or privileged material before Send?

WHY IT MATTERS

SRA Code paragraph 6.3 requires current and former clients' affairs to be kept confidential unless disclosure is required or permitted by law or the client consents. The control point matters: after-the-fact reporting cannot prevent the disclosure.

EVIDENCE TO LOOK FOR

- Pre-Send controls on the AI-chat composers the firm permits staff to use.
- Rules for privilege markers, matter and client references, personal data, credentials and firm-specific terms.
- Documented test cases, exception handling and a proportionate block, warn or redact response.

3

Does the control create another copy of the prompt?

WHY IT MATTERS

Sending a draft to a separate inspection service can introduce additional processing, retention and subprocessor questions. The firm should be able to map exactly where the message body goes for the policy decision.

EVIDENCE TO LOOK FOR

- A current data-flow diagram for the control and the AI service being used.
- A precise statement on whether message bodies or snippets are transmitted, retained or made available to support staff.
- Documented retention, access, subprocessor and data-residency arrangements.



Evidence without unnecessary collection

4

Can you prove the control acted without reading employee chats?

WHY IT MATTERS

Oversight evidence should answer what policy decision occurred, where and when, while avoiding unnecessary collection of client-confidential text. More logging is not automatically better evidence.

EVIDENCE TO LOOK FOR

- Metadata showing the site, time, decision, rule and managed user reference.
- An explicit separation between governance metadata and the message body.
- Defined access, retention, export and integrity controls for the evidence store.

5

Can you deploy, test and maintain the control as AI sites change?

WHY IT MATTERS

A control that users can disable, that fails silently, or that stops working after a site redesign is not a dependable control. Deployment and ongoing assurance matter as much as the original policy.

EVIDENCE TO LOOK FOR

- Managed browser deployment and configuration that ordinary users cannot override.
- Versioned policies, approval records, coverage tests and control-health signals.
- A process for false positives, selector changes, rule tuning and periodic review.

Quick self-assessment

Can your firm demonstrate each of these today?

- ☐ A named owner and documented AI-chat scope
 - ☐ Approved services, use cases and data classes
 - ☐ A control that acts before Send
 - ☐ No unnecessary second store of prompt content
 - ☐ Evidence that shows policy outcomes without message text
 - ☐ Managed deployment and resistance to user bypass
 - ☐ Coverage, false-positive and change-management reviews
 - ☐ Incident, exception, retention and access procedures
- 6-8 checked: review and test | 3-5: prioritise gaps | 0-2: start with one team and one workflow**

Primary sources

SRA Code of Conduct, paragraphs 6.3-6.5: sra.org.uk/solicitors/standards-regulations/code-conduct-solicitors/

Law Society, Generative AI - the essentials (updated June 2026): lawsociety.org.uk/topics/ai-and-lawtech/generative-ai-the-essentials

SRA Risk Outlook, artificial intelligence in the legal market: sra.org.uk/sra/research-publications/artificial-intelligence-legal-market/

ICO, AI security and data minimisation guidance: ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/artificial-intelligence/