



C O M P U T E R F O R E N S I C S L A B

London digital forensics & e-discovery · Established 2007

AI-Generated Documents: Provenance and Detection

Fabricated Evidence, Invented Citations and the Return to Provenance

Generative AI can now produce, in seconds, a plausible letter, contract, email chain, statement or expert-looking report, complete with confident detail and, sometimes, wholly invented facts, quotations and legal citations. This creates two distinct problems for litigation. The first is deliberate: fabricated documents introduced as genuine evidence, an email that was never sent, a contract that never existed, a reference that was never written. The second is careless: AI-assisted work product, from submissions to witness material, that contains hallucinated authorities or facts its author never checked. Detecting AI generation directly from text alone is unreliable and getting harder, so the honest forensic answer is not a magic detector but a return to first principles: provenance. Where did this document actually come from, when was it created, does it exist in the systems it should, and does its metadata and history support the account of its origin? This guide sets out for UK lawyers how AI-generated and AI-assisted documents arise in disputes, why text-based detection is weak, and how provenance, metadata and corroboration expose fabrication and protect the genuine.

🕒 Estimated reading time: 24 min read

§ ABOUT THE AUTHOR

PREPARED BY COMPUTER FORENSICS LAB E-DISCOVERY TEAM

Computer Forensics Lab is a London-based digital forensics and electronic disclosure practice established in 2007. The rise of generative AI has sharpened, rather than replaced, the laboratory's oldest discipline: establishing the provenance of a document from its metadata, its system history and its existence, or absence, across the estate. Our examiners test suspected fabricated and AI-assisted material against these objective records rather than against unreliable text detectors, and report under CPR Part 35 and CrimPR Part 19 on what a document's origin and history genuinely show.

ESTABLISHED 2007 · LONDON

ISO 17025-ALIGNED PROCEDURES

DOCUMENT PROVENANCE ANALYSIS

METADATA & AUTHENTICITY

CPR PART 35 EXPERT REPORTS

FULL CHAIN-OF-CUSTODY DOCUMENTATION

§ CONTENTS

In this guide

- 01 Executive summary
- 02 The problem in plain English: plausible text is cheap now
- 03 How AI-generated and AI-assisted documents arise
- 04 Why text-based detection is weak
- 05 Provenance: the forensic answer
- 06 Metadata, corroboration and honest limits
- 07 Deployment: fabrication, hallucinated citations and duties
- 08 Source architecture: where else the evidence lives
- 09 Worked examples
- 10 Common mistakes and technical limitations
- 11 Questions to ask · Suggested wording
- 12 Checklist and red flags · When to involve a digital forensic expert
- 13 Frequently asked questions
- 14 Glossary · References · Disclaimer · How a specialist laboratory can assist

§ 0 1 · ORIENTATION

Executive summary

The headline point: AI can now fabricate a plausible document in seconds, but detecting AI from the text alone is unreliable, so the forensic answer is not a detector but provenance: where the document actually came from, when it was created, whether it exists in the systems it should, and whether its metadata and history support its claimed origin.

How they arise (§3): deliberately fabricated documents introduced as genuine, and carelessly AI-assisted work product containing invented facts, quotations or legal citations. **Why text detection is weak (§4):** so-called AI detectors are unreliable, produce false positives and negatives, and are easily defeated, so a text-based "this was written by AI" opinion does not withstand scrutiny. **Provenance is the answer (§5):** the objective test is the document's origin and history, its metadata, authoring application, creation and revision record, and its presence or absence in the mailboxes, drives and systems where a genuine document would exist. **Metadata and corroboration (§6):** creation dates, authoring sessions and system records expose a document created after its claimed date or existing nowhere it should, corroborated across the estate. **Deployment (§7):** exposing fabricated evidence, the professional-duty problems of hallucinated citations, and the courts' expectation that authorities and facts are verified.

- **Fabrication is cheap:** AI produces plausible letters, contracts, emails and citations in seconds.
- **Text detectors are unreliable:** false positives and negatives make "written by AI" opinions unsafe.
- **Provenance is the test:** origin, creation history and existence in the right systems, not the prose.
- **Metadata exposes fabrication:** a document created after its claimed date, or existing nowhere it should, gives itself away.
- **Verify citations:** hallucinated authorities and facts are a real and rising professional-duty risk.

Two roads to authenticity: the weak one and the sound one



Figure 1 - the unreliable road is guessing AI from the text; the sound road is provenance: the document's origin, history and existence in the systems where a genuine one would live.

The problem in plain English: plausible text is cheap now

For most of legal history, producing a convincing fake document took effort, skill and time, and that friction was itself a protection. Generative AI has removed the friction. Anyone can now ask a model to produce a letter in a particular person's style, an email chain that reads like real correspondence, a contract with all the right clauses, or a reference that sounds exactly like a genuine one, and receive it, fluent and confident, in seconds. Worse, the same tools invent with equal confidence: they will supply quotations that were never said, figures that were never recorded and, notoriously, legal authorities that do not exist, all rendered as plausibly as the true material around them. The result is that plausible text can no longer be trusted on its face, whether the plausibility was manufactured deliberately to deceive or produced carelessly by someone who did not check.

The instinct is to look for an AI detector, a tool that reads a document and pronounces it machine-written. That instinct leads nowhere reliable. Text-based detection is weak, produces false positives that would wrongly brand genuine human writing as fake and false negatives that miss real fabrication, and is trivially defeated by light editing. A confident "this was written by AI" opinion, based on the prose, does not survive cross-examination. The honest forensic response is to stop asking what the text sounds like and start asking where the document came from. A genuine document has a provenance: it was created at a certain time, in a certain application, by a certain account, and it exists in the mailboxes, drives and systems where a real document of its kind would be found. A fabricated one usually does not, and its metadata, its creation history and its absence from the systems it should inhabit betray it far more reliably than any analysis of its words. This guide is about that return to provenance.

How AI-generated and AI-assisted documents arise

- **Deliberate fabrication:** a document generated to deceive and introduced as genuine, an email never sent, a contract never signed, a letter never written, indistinguishable on its face from a real one: the fraud the return to provenance is designed to expose (§5).
- **Fabricated supporting detail:** genuine documents padded or altered with AI-generated passages, invented quotations or figures woven into otherwise real material: partial fabrication, exposed by revision history and internal inconsistency (§6).
- **Hallucinated citations and facts:** AI-assisted submissions, statements or reports containing legal authorities, cases and facts that do not exist, produced carelessly rather than dishonestly, but no less damaging: a professional-duty problem as much as an evidential one (§7).
- **Synthetic correspondence and records:** whole email chains, chat logs or account records fabricated to support a narrative, which must be tested against the systems that would independently hold genuine versions (§8): the fake defeated by the absence of a real counterpart.
- **The honest and disclosed use:** legitimate AI-assisted drafting, properly checked and, where required, disclosed, distinguished from the fabrication and the negligent, so the tool itself is not treated as the wrong (§7): the line between assistance and deception.

§ 0 4 · WHY TEXT DETECTION IS WEAK

Why text-based detection is weak

- **Detectors are unreliable:** AI-text detectors produce significant false positives, wrongly flagging genuine human writing, and false negatives, missing real AI output, so their verdicts are unsafe as evidence: the tool that cannot be trusted either way.
- **Easily defeated:** light editing, paraphrase or running text through another tool defeats detection, so even a working detector is evaded with trivial effort: the arms race the detector loses.
- **No ground truth in the text:** fluent, confident prose is now producible by human and machine alike, and there is no reliable signature in the words themselves that distinguishes them: the reason a text-only opinion lacks foundation.
- **The technology moves:** models change constantly, so any detector is chasing a moving target and its accuracy degrades: an unstable basis for an expert opinion (guide 100's honest-limits discipline).
- **The wrong question:** "does this read like AI?" is the wrong question; "where did this document come from and does its history support its claimed origin?" is the answerable one, which is why the analysis moves to provenance (§5).

Provenance: the forensic answer

- **Origin and authoring:** the document's own metadata, the authoring application, author identifiers, creation and last-modified records and, for Office files, the internal revision and editing history (guide 4's document-metadata discipline), tested for consistency with the claimed origin.
- **Creation timing against the claim:** the objective creation and modification records compared with the date the document purports to bear, exposing a document created long after its claimed date, the classic backdating tell (guide 109's world-test discipline).
- **Existence in the right systems:** the decisive test, whether the document exists where a genuine one would, in the sender's sent items and the recipient's inbox, on the server, in the backups, in the version history, because a real email or file leaves corroborating copies and a fabricated one usually does not (§8): presence, or telling absence.
- **Internal and cross-document consistency:** fonts, formatting, template artefacts, embedded objects and metadata checked for internal consistency and against genuine documents of the same kind, since fabrications often carry anachronisms and mismatches (guide 4).
- **The chain of custody of the original:** the native original obtained and examined, not a printout or screenshot, because provenance lives in the native file and its system context, and a flattened copy strips the very metadata the analysis depends on (guides 2, 3).

Metadata, corroboration and honest limits

- **Metadata can be edited too:** metadata is powerful but not infallible, it can itself be altered, so it is corroborated against independent system records rather than taken alone, the same discipline as any artefact (guide 108).
- **Absence is strong, presence is expected:** the absence of a supposedly genuine document from every system that should hold it is powerful evidence of fabrication; its presence, consistent across systems, supports authenticity: the corroboration cutting both ways (§5).
- **AI use is not proven from text:** the honest limit, that an examiner establishes fabrication and false provenance, not "this was generated by a particular AI", which the text cannot reliably show (§4): the claim confined to what the evidence supports.
- **Genuine documents have innocent quirks:** a metadata oddity may have an innocent explanation (a template, a conversion, a clock error), so anomalies are investigated and contextualised, not assumed to prove fabrication (guide 100).
- **State the finding at its level:** the conclusion expressed as what provenance shows, "this document was created after its claimed date and exists in no system that should hold it", rather than an unsupportable "this is AI-generated": the honest, survivable finding.

Deployment: fabrication, hallucinated citations and duties

- **Exposing fabricated evidence:** a document proved by provenance to be created after its claimed date, or to exist nowhere it should, put to the court as fabricated, with the objective records rather than a text opinion carrying the finding (§5, guide 98).
- **Hallucinated authorities:** the growing problem of AI-assisted submissions citing cases and authorities that do not exist, exposed by simple verification, and engaging serious professional-conduct and duty-to-the-court concerns for the lawyer responsible (§3): a matter courts increasingly police.
- **Verifying before relying:** the practical duty to check AI-assisted work, every authority, quotation and fact verified against a primary source before it is put before a court or opponent, because the tool invents confidently (§4).
- **Protecting the genuine:** the defensive use, provenance establishing that a genuine document challenged as an AI fabrication is authentic, because it exists consistently across the systems and its history supports its origin (§5): the innocent document defended.
- **Disclosure and honest use:** the legitimate, checked and, where appropriate, disclosed use of AI in preparing material distinguished from fabrication and negligence, so the profession's honest adoption of the tools is not conflated with their abuse (guide 20's proportionality, and current Law Society and court guidance).

Source architecture: where else the evidence lives

Per this series' source-architecture discipline: a document is never only itself, and its authenticity is decided by whether the systems that should corroborate it actually do. The document holds its own metadata and revision history. But a genuine email exists in the sender's sent items and the recipient's inbox and on the mail server and in the backups; a genuine file exists on the drive where it was made, in its version history and in the cloud sync; a genuine record exists in the application or database that would have produced it; and the counterparties hold their independent copies. The discipline is to test the suspect document against every system that a real one of its kind would inhabit, because fabrication is exposed not by the document's words but by its absence from the places it should be and the inconsistency of the few places it appears. When a fabricated email exists only as a single file with no counterpart in any mailbox or server, the absence condemns it; when a backdated contract's creation metadata and the drive's own records contradict its date, the history condemns it; and when a genuine document is challenged, its consistent presence across every corroborating system defends it.

CORROBORATING SOURCE	THE DOCUMENT ITSELF	SENDER / AUTHOR SYSTEM	RECIPIENT / COUNTERPART	SERVER / CLOUD	BACKUPS / VERSIONS	APPLICATION / DATABASE
Email	Metadata, headers	Y: sent items	Y: inbox copy	Y: mail server	Y: mail backups	n/a
Office file	Author, revision history	Y: drive of origin	Shared copies	Y: cloud sync	Y: version history	n/a
Contract / letter	Creation metadata	Y: authoring record	Y: other party's copy	Cloud copy	Prior versions	DMS record
System record	Export metadata	n/a	Counterpart export	Server data	DB backups	Y: source database
Creation date	Claimed on face	Y: true create time	Received time	Server timestamps	Backup dates	Log timestamps

Fallback strategy in one line: test the document against every system a genuine one would inhabit; fabrication betrays itself by absence and inconsistency, and authenticity is proved by consistent presence across the corroborating sources.

Worked examples

EXAMPLE 1 · THE EMAIL THAT EXISTED ONLY AS A FILE

A party produced a decisive email, apparently sent two years earlier, that supported its case perfectly. It was, on its face, entirely plausible, the kind of fluent, well-judged message AI now produces effortlessly. The §5 provenance test, not any reading of its text, exposed it. The email existed only as a single file the party had produced; it appeared in no sent-items folder of the supposed sender, in no inbox of the supposed recipient, on no mail server and in no backup, precisely the §8 absence that a genuine email of that age could not exhibit. Its internal headers, examined natively, were inconsistent with the mail systems it claimed to have passed through, and its file-creation metadata post-dated its purported send by two years (§6). No text detector was needed or used; the objective records showed a document that existed nowhere a real one would. The lesson is §5's: a fabricated document is exposed by provenance, its absence from every system that should hold it and the contradiction of its own metadata, not by an unreliable opinion on its prose.

EXAMPLE 2 · THE SUBMISSION WITH AUTHORITIES THAT DID NOT EXIST

A set of submissions cited several authorities in support of a proposition, confidently and in proper form. Routine verification, the §7 duty to check, revealed that two of the cited cases did not exist: they had been produced by an AI tool used in drafting and never verified against a primary source, a §3 hallucination rather than a deliberate deceit. The problem was not detected by analysing whether the text "sounded like AI", it was found by looking the citations up, the simplest and most reliable check of all. The episode engaged the drafting lawyer's duty to the court and professional-conduct obligations, and illustrated that the real defence against hallucinated material is verification, not detection. The lesson is §7's: AI invents authorities and facts with complete confidence, and the only reliable safeguard is to verify every citation, quotation and fact against a primary source before it is relied on, a duty the courts increasingly enforce.

EXAMPLE 3 · THE GENUINE DOCUMENT DEFENDED BY ITS PROVENANCE

A party challenged a genuine internal report as an AI-generated fabrication, pointing to its fluent style. The §5 provenance analysis defended it precisely because the document had a real history: it existed on the author's drive with a consistent creation and revision record spanning weeks (§6), earlier draft versions were present in the version history and cloud sync, it had been emailed to colleagues who held it in their mailboxes, and every corroborating system in the §8 estate contained a consistent copy. Against that objective record, the stylistic challenge collapsed: whatever the report read like, its provenance was that of an ordinary, genuinely-authored document. The lesson is that provenance protects the innocent as well as exposing the guilty, a genuine document with a consistent creation history and corroborating copies across the systems is authenticated by its origin, and a challenge based only on how the text sounds cannot overcome it.

Common mistakes and technical limitations

Common mistakes

- **Relying on an AI detector:** the fundamental error, resting a fabrication case on an unreliable text detector rather than provenance (Example 1, §4).
- **Working from a printout or screenshot:** the native original never obtained, so the metadata and system context that decide provenance are lost (§5).
- **Not checking the systems:** the document assessed in isolation, without testing whether it exists where a genuine one would (Example 1, §8).
- **Failing to verify citations:** AI-assisted authorities and facts relied on unchecked (Example 2, §7).
- **Taking metadata at face value:** metadata trusted without corroboration, when it too can be edited (§6).
- **Assuming fluent means fake:** a genuine document branded a fabrication because it reads well (Example 3).
- **Conflating assistance with deception:** honest, checked AI-assisted drafting treated as if it were fabrication (§7).
- **Overclaiming the source:** asserting a document was made by a specific AI, which the evidence cannot show (§6).

Technical limitations

- **Text detection does not work reliably:** false positives and negatives make it unsafe, and it is easily defeated: the central limit (§4).
- **Metadata is not infallible:** it can be altered, so it is corroborated, not taken alone (§6).
- **Absence needs the right systems:** the power of "it exists nowhere" depends on having access to the systems that should hold it (§8).
- **Provenance needs the native file:** flattened copies strip the metadata the analysis requires (§5).
- **The finding is fabrication, not authorship:** provenance shows false origin, not which tool produced the text (§6).

§ 11 · INTERROGATORIES & DRAFTING AIDS

Questions to ask · Suggested wording

Ask your client

- For any key document, do we have the native original with its metadata, and can we obtain it from the system where it was created?
- Where should a genuine version of this document also exist, mailboxes, drives, the other party, backups, so we can test its provenance?
- Was any AI tool used in preparing material in this matter, and has every authority, quotation and fact been verified?

Ask your opponent

- Please produce the native electronic original, with metadata intact, of [document], and identify every system in which a genuine copy exists (sender and recipient mailboxes, servers, drives, backups, version history).
- Please confirm the creation, authoring and revision history of [document] and the account and application that created it.
- Please confirm whether AI tools were used in preparing any document relied on, and whether the authorities and facts in it have been verified against primary sources.

Ask your eDiscovery / forensic provider

- Does this document's metadata and creation history support its claimed date and origin?
- Does it exist, consistently, in the systems where a genuine document of its kind would live?
- What can, and cannot, be honestly said, fabrication and false provenance yes, "written by AI" no?

SUGGESTED WORDING · INSTRUCTION FOR A DOCUMENT-PROVENANCE EXAMINATION

"We instruct you to examine the provenance and authenticity of the document(s) at Schedule 1, which are suspected of being fabricated or AI-generated. Please: (1) obtain and examine the native electronic original with its metadata intact, not a printout or screenshot, and preserve integrity throughout; (2) analyse the authoring application, author identifiers, creation and modification records and any internal revision or editing history, and compare the objective creation timing against the date the document purports to bear; (3) test whether the document exists, consistently, in every system where a genuine document of its kind would be found (sender and recipient mailboxes, mail servers, drives, cloud sync, version history, backups and any source application or database), treating absence and inconsistency as evidence of fabrication and consistent presence as support for authenticity; (4) do not rely on text-based AI-detection tools, and confine any conclusion to what provenance shows, fabrication and false or inconsistent origin, rather than an unsupportable assertion that the text was written by AI; and (5) investigate and contextualise any metadata anomaly for innocent explanations before treating it as proof."

Checklist and red flags · When to involve a digital forensic expert

The provenance checklist

- ☐ Native original obtained, with metadata intact, not a printout
- ☐ Authoring application, author and creation records examined
- ☐ Creation timing compared against the claimed date
- ☐ Internal revision and editing history reviewed
- ☐ Existence tested across every system a genuine copy would inhabit
- ☐ Absence and inconsistency weighed as fabrication evidence
- ☐ Metadata corroborated, not taken at face value
- ☐ Text-based AI detectors not relied on
- ☐ Citations, quotations and facts verified against primary sources
- ☐ Findings confined to provenance, not "written by AI"

Red flags

- A decisive document existing only as a produced file, with no counterpart anywhere: Example 1.
- Cited authorities that cannot be found in any primary source: Example 2.
- Only a printout or screenshot available, no native original.
- Creation metadata post-dating the document's claimed date.
- A fabrication case resting on a text "AI detector".
- A genuine document challenged solely because it reads fluently.
- Metadata relied on with no corroboration.

When to involve a digital forensic expert

Whenever a document's authenticity is genuinely in issue in the age of easy fabrication, because the sound method is provenance and it is a forensic discipline: the expert obtains the native original, analyses its metadata and creation history, and, decisively, tests whether it exists where a genuine document would across the estate, so fabrication is shown by objective records rather than an unreliable text opinion (Example 1, §5, §8); on the citation problem, to support the verification of AI-assisted material against primary sources (Example 2, §7); in defence, to authenticate a genuine document challenged as a fabrication by proving its consistent provenance (Example 3); and at deployment, where the finding is presented under CPR Part 35 or CrimPR Part 19, confined honestly to what provenance shows and not overstated as "written by AI". Through **cflab.uk** and **e-discovery.uk**, the response to AI-generated documents is the oldest discipline made newly essential: not guessing at the text, but establishing where a document truly came from, when, and whether the systems that should corroborate it actually do.

§ 1 3 · COMMON QUESTIONS

Frequently asked questions

Can you tell if a document was written by AI?

Not reliably from the text, and no honest examiner will claim to (§4). AI-text detectors produce false positives and negatives and are easily defeated, so a "this was written by AI" opinion based on the prose does not survive scrutiny. What can be established, objectively, is provenance: whether the document's metadata and creation history support its claimed origin, and whether it exists where a genuine document would (§5). That, not text analysis, is the reliable answer.

How do you prove a document is fabricated?

By provenance and corroboration, not prose (§5, Example 1). A genuine email exists in mailboxes, on servers and in backups; a genuine file has a consistent creation and revision history and copies in the cloud and with others. A fabricated document typically exists only where the party produced it, contradicts its own creation metadata, and is absent from every system that should hold it (§8). That objective absence and inconsistency, not how the text reads, exposes the fake.

Why do we need the native file and not a printout?

Because provenance lives in the native file and its system context (§5). A printout or screenshot strips the metadata, the authoring records, the creation and revision history, the very things the analysis depends on, and flattens the document into an image that cannot be tested. To examine provenance you need the electronic original with its metadata intact, obtained from the system where it lives, with integrity preserved.

Our opponent's submissions cite cases we cannot find. What does that mean?

It may mean the authorities were hallucinated by an AI tool used in drafting and never verified (§3, Example 2), an increasingly common and serious problem. The check is the simplest one: verify each citation against a primary source. Authorities that exist nowhere are a real risk to the party relying on them, engaging duties to the court, and the safeguard is verification, not any attempt to "detect AI" in the writing.

Can metadata itself be faked?

Yes, metadata can be altered, which is why it is never taken alone (§6). Its power comes from corroboration: a creation date is tested against independent system records, and a document's claimed metadata is checked against the mailboxes, servers and backups that should independently reflect it. Consistent corroboration across systems supports authenticity; contradiction and absence expose fabrication. Metadata is a strong tool, used with the same corroboration discipline as any artefact.

Is using AI to draft documents a problem in itself?

No, honest, checked and, where required, disclosed AI-assisted drafting is legitimate, and this guide does not treat the tool as the wrong (§7). The problems are fabrication, using AI to manufacture false evidence, and negligence, relying on invented authorities or facts without verifying them. The line is between assistance that is checked and disclosed and deception or carelessness that is not, and provenance and verification, not hostility to the technology, are how that line is policed.

§ 1 4 · REFERENCE

Glossary

Generative AI

Tools that produce fluent text, including plausible fabrications.

Hallucination

Confident AI output that is invented, such as non-existent cases.

Provenance

Where a document came from, when, and by what path.

Native original

The electronic file with metadata intact, not a printout.

Metadata

Data about a document: author, creation and revision records.

Revision history

The internal record of a document's edits over time.

AI detector

A text tool guessing at AI authorship; unreliable as evidence.

Corroboration

Confirmation from independent systems that should hold a copy.

Backdating

Giving a document a false earlier date, exposed by creation records.

Authentication

Establishing a document is what it purports to be.

Sources and authoritative references

The provenance and authenticity disciplines in this guide draw on materials published by our laboratory and e-discovery practice, referenced here as sources:

- cflab.uk, digital forensics services (document provenance, metadata and authenticity analysis, CPR Part 35 and CrimPR Part 19 reporting): cflab.uk/digital-forensics-services · guides library: cflab.uk/guides
- e-discovery.uk, EDRM Phase 06 · Analysis and Phase 07 · Review (authenticity testing and AI-assisted review as practised): e-discovery.uk/edrm/analysis · e-discovery.uk/edrm/review
- e-discovery.uk, practice method and Knowledge Centre: e-discovery.uk · e-discovery.uk/knowledge
- Law Society, guidance on generative AI and on its use in legal practice and disclosure: lawsociety.org.uk, [generative AI](#) · Courts and Tribunals Judiciary, guidance on AI for judicial office holders: judiciary.uk, [AI guidance](#)
- Forensic Science Regulator, Code of Practice v2 (authenticity and interpretation): gov.uk, [FSR Code](#) · CPR Part 35 (expert evidence and duties): justice.gov.uk, [Part 35](#)
- PD 57AD and CPR Part 31: justice.gov.uk, [PD 57AD](#) · justice.gov.uk, [Part 31](#)

DISCLAIMER

This publication provides general information about AI-generated documents, provenance and detection as at August 2026. AI tools and their capabilities change rapidly; text-based AI detection is unreliable and this guide does not endorse it, and conclusions about fabrication should rest on provenance and corroboration. The worked examples are fictional. This guide is not legal advice, does not create a professional relationship, and should not be relied upon in place of advice on the facts of a specific matter from a qualified lawyer or a suitably qualified forensic practitioner. Links were live at the date of publication.

§ HOW A SPECIALIST LABORATORY CAN ASSIST

Working with Computer Forensics Lab

The laboratory's response to AI-generated documents is deliberately unfashionable: it does not reach for a detector, because detectors do not work, but for provenance, the discipline that decided document authenticity long before generative AI and decides it still. The native original is obtained with its metadata intact, its authoring application, author, creation and revision history are examined, and its creation timing is tested against the date it claims (§5, §6). The decisive step is corroboration across the estate: whether the document exists, consistently, in every system a genuine one of its kind would inhabit, the mailboxes, servers, drives, version history and backups, so fabrication is exposed by objective absence and inconsistency and authenticity is proved by consistent presence (Examples 1 and 3, §8). The same discipline supports verification of AI-assisted material, where the safeguard against hallucinated authorities is checking them against primary sources, not analysing prose (Example 2). Every finding is confined honestly to what provenance shows, fabrication and false origin, never an unsupportable "written by AI". Reports run under CPR Part 35 or CrimPR Part 19. Conflicts checks, confidential scoping discussions and fixed-fee estimates are routine; and the principle is constant: in an age when plausible text is cheap, the question that matters is where the document truly came from.



INSTRUCT THE LAB

Speak to a forensic examiner, not a salesperson.

Describe your matter in confidence and an examiner will respond the same working day. Conflicts checks and scoping calls with US counsel are routine.

NEW ENQUIRIES

+44 (0)20 7164 6971

EMAIL

info@cflab.uk

E-DISCOVERY

e-discovery.uk

Computer Forensics Lab Ltd · Euro House, 133 Ballards Lane, Finchley, London N3 1LJ, United Kingdom

cflab.uk · Customer service +44 (0)20 7164 6915 · ICO ZB395023 · NPCC · ISO 17025 aligned practice · UK Gov Digital Marketplace