COMPUTER FORENSICS LAB

London digital forensics & e-discovery · Established 2007

Forensic Collection from Corporate File Servers

Shared Drives, Permissions, Snapshots, Deleted Files and the Attribution Problem

The shared drive is where corporate disputes live: the contracts folder, the finance share, the project tree that everyone touched and nobody owns. Collecting it forensically differs from imaging a laptop at every step: servers cannot be switched off and carried away, permissions and audit records are part of the evidence, snapshots hold the history, deletion behaves differently, and the question "who did this?" runs through accounts and endpoints rather than a single user's machine. This guide explains how file-server evidence is preserved and collected defensibly, what the server records about access and change, how deleted and historical material is recovered, and how attribution on a shared system is actually established.

⏱ Estimated reading time: 27 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. Server-side collection is standing work for the laboratory: live corporate estates, snapshot-based acquisition, permission and audit analysis, and the attribution work shared systems demand, delivered under the protocol, custody and reporting disciplines of this series. Its eDiscovery arm, **e-discovery.uk**, processes server collections at disclosure scale.

ESTABLISHED 2007 · LONDON | ISO 17025-ALIGNED PROCEDURES | ISO 27001-ALIGNED SECURITY | ACPO / NPCC DIGITAL EVIDENCE PRINCIPLES | SERVER & LIVE-SYSTEM ACQUISITION | CPR PART 35 EXPERT REPORTS | FULL CHAIN-OF-CUSTODY DOCUMENTATION

§ CONTENTS

In this guide

§ 0 1 · ORIENTATION

Executive summary

THE HEADLINE POINT: COLLECT THE FILES, THE HISTORY AND THE CONTEXT, OR THE COLLECTION ANSWERS NOTHING

A corporate file server holds three layers of evidence, and disputes usually need all three. The **content layer** is the files themselves, collected forensically (metadata preserved, hashed, manifested) from the live system or, better where available, from a storage snapshot that freezes the estate at a defensible moment. The **history layer** is what the server remembers about the past: snapshot series and previous-versions data holding earlier states and deleted files, backup generations extending the reach to years, and file-system internals recording creation, modification and (sometimes) deletion events. The **context layer** is what makes shared-drive evidence usable: the permission structure (who could reach which folders, and how that changed), the audit logs (who actually opened, changed, renamed or deleted what, when, from which endpoint, where auditing was on), and the mapping from accounts and endpoints to people. Collections that take only the content layer produce a pile of documents that cannot answer the questions that made anyone collect them: who put this here, who changed it, who took it, who could even see it. The lawyer's role is to make sure the instruction, the protocol and the preservation notice name all three layers, early, because the history layer rotates (snapshots prune, logs overwrite, backups recycle) on schedules that do not wait for proceedings.

- **Servers are collected live or by snapshot, not carried away:** the machinery of guide 33's live systems: imaging a production server to a standstill is rarely proportionate or necessary; snapshot-based and share-level forensic collection, logged and hashed, is the standard and is defensible when done to protocol.
- **Timestamps need server-side literacy:** created/modified dates behave differently across copies, migrations and restores; a share migrated last year can make every document "created" on migration day, and the examiner reads dates with that history in hand or misreads everything.
- **Audit logging is conditional:** file-access auditing exists but is often partially enabled or off; what was configured, since when, is a first-day question, because it defines whether "who opened it" is answerable from logs or must be built from endpoint artefacts instead.
- **Deletion is shallower on servers, in both directions:** ordinary users' deletions often bypass any recycle mechanism on shares, but snapshots and backups usually catch what user-side deletion destroys: the recovery conversation is about retention schedules, urgently.
- **Preserve the schedules themselves:** the single most valuable early act is suspending snapshot pruning, log rotation and backup recycling for the relevant systems: one email to IT, drafted in §10, buying months of history that would otherwise expire routinely.

§ 0 2 · FIRST PRINCIPLES

The problem in plain English: everyone's drive, no one's machine

When a laptop matters, the questions are about one person's world. When the finance share matters, the questions are about a crowd's: forty people could reach the folder; the disputed spreadsheet has been edited by six accounts over three years; the version that matters was deleted, or replaced, or renamed, at some point nobody can name; and the drive itself has been migrated twice, restored once and reorganised by a contractor. The evidence is richer than any laptop's, and the naive reading of it is more dangerous than any laptop's, because every artefact (a date, a name in the "modified by" column, a file's presence in a folder) is one operation removed from what it seems to say.

The discipline that tames this is layered collection with server literacy: take the content forensically; take the history (snapshots, versions, backups) before the schedules eat it; take the context (permissions, audit configuration and logs, the account-to-person map); and interpret dates and attributions through the system's actual biography: migrations, restores, sync clients and admin activity included. Done this way, the shared drive becomes what it really is: the best-documented crime scene in the building. Done naively, it becomes the source of confident findings that dissolve under one competent question: "and how do you know a person, rather than the backup agent, touched that file?"

§ 0 3 · THE DIFFERENCES THAT MATTER

How server collection differs from device imaging

- **Availability rules the method:** production servers serve the business; the collection options, in rough order of preference where offered by the estate, are: storage or VM snapshots (freeze the state, collect from the frozen copy at leisure); share-level forensic collection from the live system (metadata-preserving tools, out-of-hours windows); and full imaging only where the server itself is contested or small enough to make it proportionate: the guide 32 decision tree applied to iron that cannot stop.
- **Scale rules the scope:** shares run to terabytes and decades; guide 37's targeting (named paths, custodial folders, date bounds, with the preserve-broad backstop via snapshots) is not optional economics but the difference between a collection that completes and one that occupies a quarter.
- **The system is multi-layered:** Windows file servers, clustered storage, virtualised guests on hypervisors, and NAS appliances (the next guide) each put the evidence in different places; the reconnaissance of guide 39 establishes which world this share lives in before the method is chosen.
- **Admin cooperation is structural:** collection runs through administrative access: snapshots, permissions exports, audit retrieval: so the roles-and-supervision disciplines (IT assists, examiner executes, everything logged) carry the defensibility, and the admin's own account activity becomes part of the record.
- **The evidence includes configuration:** what was collected is only interpretable alongside how the system was set up: permission model, audit policy, snapshot and backup schedules, quota and archive behaviour: exported and exhibited as part of the collection, not reconstructed from memory a year later.

§ 0 4 · WHAT THE SERVER KNOWS

What the server records: permissions, audit and metadata

- **Permissions are a map of possibility:** the folder-by-folder access-control structure (users, groups, inheritance) answers who *could* see or change material: probative in confidentiality disputes ("the pricing folder was open to the whole company" cuts both ways), and in narrowing suspects. Group membership histories and permission-change events (where audited) add the time dimension: who was given access, when, by whom.
- **Audit logs are a record of actuality, where enabled:** Windows file-access auditing (and equivalents on other platforms) can log opens, writes, renames, permission changes and deletions per account with timestamps and source endpoints. The two first-day facts: what auditing was configured (often only some events, some folders), and the log retention (security logs overwrite; weeks is common). Where auditing was rich, the share tells its own story; where it was off, actuality is rebuilt from endpoint artefacts (the Windows and Mac guides) and network traces.
- **File metadata carries operations history, encoded:** owner fields, created/modified stamps and (on NTFS) file-system internals record events, but through the lens of every copy, migration, restore and sync since: the §2 literacy point. Embedded document metadata (authors, revision histories: this series' document-forensics territory) rides along and often outlives the file-system story.
- **Access is also recorded elsewhere:** authentication logs (who connected to the server, from which machine, when), VPN records for remote access, and endpoint-side artefacts (mapped-drive usage, LNK files and shellbags pointing at server paths, per guide 41) triangulate the server's own account: the multi-source convergence this series keeps returning to.

§ 0 5 · THE HISTORY LAYER

Snapshots, previous versions and deleted files

- **Snapshot series are the share's memory:** enterprise storage and Windows' own machinery (Volume Shadow Copies exposed as "Previous Versions") keep periodic point-in-time states: hourly, daily, weekly, on retention schedules: each holding files as they then stood, deleted files included. For any "what did the folder look like on the 14th?" question, the snapshot nearest the 14th is the answer, mounted and collected forensically.
- **Backups extend the series:** nightly/weekly/monthly generations reach months and years back where snapshots reach days and weeks; the dedicated backup guide later in this series covers the search-obligation questions, but the preservation point belongs here: recycling schedules eat generations on a timetable, and the §10 suspension email stops the clock.
- **User deletion on shares is often unforgiving forward, forgiving backward:** deleting over the network typically skips client-side recycle bins, so the live share loses the file at once: but the snapshot from an hour earlier holds it, which is why the deletion cases on servers are usually won from the history layer, not from carving.
- **Deletion events themselves:** where auditing captured deletes, the log names the account, time and endpoint; where it did not, deletion windows are bracketed by snapshot differencing (present in Monday 02:00, absent in Tuesday 02:00) and correlated with authentication and endpoint evidence: narrower than a log line, regularly sufficient.
- **Snapshot differencing at scale is a method in itself:** comparing successive states yields the change history of a whole tree (what appeared, changed, vanished, per interval): the machine-generated chronology of a project folder that no witness could write, and a standard exhibit in our server matters.

§ 0 6 · WHOSE HAND

Attribution on shared systems

- **The chain has four links:** file event → account → session/endpoint → person. Server evidence (audit, metadata, authentication) usually supplies the first two; endpoint and access evidence (guide 41's session work, VPN and badge records) supplies the third; and the human layer is argued from credential exclusivity, patterns and elimination, as ever. Reports that jump from "modified by jsmith" to "Mr Smith altered it" have skipped two links, and get taken apart accordingly.
- **Service and admin accounts muddy the water knowingly:** backup agents, sync services, migration tools and administrators touch files at scale, updating metadata and appearing in logs; the examiner separates mechanical from human activity by account type, pattern and timing, and the honest chronology labels which is which.
- **Shared credentials are the recurring wound:** the departmental login everyone knows converts the account link into a suspect pool; attribution then leans entirely on the endpoint and human layers, and the report says so plainly: this series' shared-computer example, at share scale.
- **Endpoint correlation is the closer:** the server says account X deleted the folder at 17:12 from endpoint WS-014; WS-014's own artefacts (logon session, shellbags showing the share path, the VPN record if remote) put a person in the chair: the two-sided collection (server plus implicated endpoints) that the protocol should anticipate from the start.

§ 0 7 · THE PLAYBOOK

The collection playbook

- **Day one, preserve the perishable:** the §10 email: suspend snapshot pruning, log rotation and backup recycling for named systems; freeze permission changes on disputed trees where operationally possible; export current audit configuration; and note the system biography (migrations, restores, reorganisations) while IT remembers it.
- **Scope from the map:** identify the shares, paths and date ranges in issue (guide 37's criteria, from interviews and the data map), and the snapshot/backup points that bracket the events: the collection plan names states as well as paths.
- **Collect by protocol:** guide 36's document, with the server specifics: snapshot creation or selection as the acquisition moment; metadata-preserving share collection with per-item hashing and manifests; permission and audit exports as first-class exhibits; and the deviation log for the live estate's inevitable frictions.
- **Verify and reconcile:** item counts against source, sample re-checks, and the manifest that lets "complete within criteria" be arithmetic (guide 37's discipline, unchanged at server scale).
- **Deliver all three layers together:** content set, history set (snapshot/backup extracts and differencing), context set (permissions, audit, configuration, account map), one custody record: the package from which both review and expert analysis can proceed without re-visiting the estate.

§ 0 8 · IN THE WILD

Worked examples

EXAMPLE 1 · THE VANISHED CONTRACTS FOLDER, REBUILT FROM SNAPSHOTS

Mid-dispute, a supplier's "Contracts/2021-2023" folder was found empty, its counterparty crying destruction. The storage system's snapshot series told a calmer and stranger story: differencing showed the tree intact until a Thursday 21:40 window, then gone; the audit log (deletes were audited) attributed the deletion to a named admin account from the IT room's workstation; and the same evening's tickets showed a storage-space purge of "old folders" by a contractor working from a list that should not have included the tree. The folder was restored in full from the prior night's snapshot, hashes matching earlier backup generations; the destruction allegation collapsed into a negligence footnote; and the case proceeded on documents that, without the history layer, would have been "destroyed" in every pleading thereafter.

EXAMPLE 2 · "MODIFIED BY" AND THE MIGRATION THAT FOOLED EVERYONE

A claimant's schedule listed forty documents "modified by" the defendant in a single week, presented as a falsification spree. The server's biography dissolved it: that week was the share's migration to new storage, performed under an admin account that the migration tool impersonated per-file; every document on the share carried the same week's stamps; and the defendant's account appeared because the tool preserved original ownership fields inconsistently. Snapshot states from before the migration showed the documents unchanged for years. The expert report's migration section, with the IT change records exhibited, removed the spree from the case in two pages: server dates, read without the system's history, had nearly built a fraud claim out of routine maintenance.

EXAMPLE 3 · THE PRICING FOLDER FORTY PEOPLE COULD READ

An employer alleged its departed salesman had accessed "highly restricted" pricing models before joining a rival. The permission export, taken as part of collection, showed the folder's access list: a group containing the entire commercial department, forty-one accounts, unchanged for three years: and the audit configuration showed file-open events were never logged for that tree. The claim's "restricted" limb failed on the map of possibility alone; the actuality limb, rebuilt honestly from his endpoint (guide 41's recency artefacts showed two openings, both within his role's ordinary pattern), supported no adverse inference. The case settled on nominal terms, and the permission export (one command, ten minutes) had been the decisive exhibit.

§ 0 9 · WHERE IT GOES WRONG

Common mistakes and technical limitations

Common mistakes

- **Content-only collection:** the files taken, the history and context layers left to rot on their schedules: §1's central failure.
- **IT copies the share:** drag-and-drop by administrators, timestamps rewritten, per this series' founding anti-pattern, at terabyte scale.
- **Preservation email never sent:** snapshots pruning and logs rotating through the letter-before-action period; months of history lost to politeness.
- **Dates read without the biography:** Example 2's migration fallacy; restores and sync clients produce it equally.
- **Attribution from the "modified by" column:** the four-link chain collapsed into one word.
- **Service-account noise reported as human activity:** the backup agent indicted for touching every file nightly.
- **Snapshots examined but not preserved:** conclusions drawn from states that then pruned away, leaving findings without exhibits.
- **The admin unexamined:** collections that never export who held administrative power, in disputes where that is the question.

Technical limitations

- **Audit answers only what was asked:** auditing off, partial or short-retained leaves actuality to be rebuilt from endpoints; the configuration export defines the ceiling honestly.
- **Snapshots are samples, not video:** events between snapshots are bracketed, not timestamped; differencing gives intervals, and the report says intervals.
- **Carving rarely rescues shares:** enterprise storage abstractions (RAID, thin provisioning, dedupe) make unallocated-space work on servers a specialist, often low-yield exercise; the history layer is the recovery route, which is why its preservation dominates everything.
- **Scale constrains completeness claims:** terabyte trees are collected within criteria and verified by manifest and sample, per guide 37; "we took everything ever on the share" is rarely true and never needs to be.
- **Cloud-hosted shares change the rules:** the same folder served from SharePoint or a cloud file service has different history and audit machinery: the cloud guides ahead; establish which world the "shared drive" actually inhabits before instructing.

§ 10 · INTERROGATORIES & DRAFTING AIDS

Questions to ask · Suggested wording

Ask your client

- What is the share's biography: platform, migrations, restores, reorganisations, sync clients? (The date-literacy layer.)
- What snapshot, backup and log retention schedules apply, and who can suspend them today?
- What file auditing is configured, since when, and where do those logs live?

Ask your opponent

- Please confirm suspension of snapshot pruning, log rotation and backup recycling for [systems] pending collection, and state the retention schedules that applied before suspension.
- Please provide the permission structure and audit configuration for [paths] as at [date], with change histories where held.
- Please identify all administrative and service accounts with write access to [paths] in the period, and the persons controlling each.

Ask your eDiscovery / forensic provider

- What acquisition route fits this estate (snapshot, live share collection, imaging), and why?
- How will the three layers be delivered: content manifest, snapshot differencing, permission and audit exports?
- What does the attribution plan look like: which endpoints will need collecting alongside the server?

SUGGESTED WORDING · DAY-ONE PRESERVATION EMAIL TO IT (ADAPT AND SEND TODAY)

> "With immediate effect and until further notice, for [named servers/shares/systems]: (1) suspend all automatic deletion, pruning and recycling of storage snapshots, previous-versions data, backup generations and event or audit logs, extending retention as necessary; (2) make no permission changes, reorganisations, migrations or clean-ups affecting [paths] without our written confirmation; (3) export and preserve the current permission structure and audit configuration for [paths]; and (4) record and report to us any scheduled maintenance that cannot be deferred. Please confirm each step and copy this instruction to any managed-service provider involved."

SUGGESTED WORDING · INSTRUCTION FOR A SERVER EXAMINATION

> "Please collect from [system] the content of [paths] as at [date/snapshot], with metadata preserved, hashed and manifested; the snapshot and backup states bracketing [events], with a differencing analysis of changes to [paths] between [dates]; and the permission structure, audit configuration and available audit logs for the period. Please report on: the creation, modification, deletion and movement of [categories], with each event attributed through the account, endpoint and (so far as the evidence allows) human layers, distinguishing service and administrative activity; and any effect of migrations, restores or synchronisation on the metadata relied on. Please state audit coverage and snapshot intervals as limitations where applicable."

§ 11 · QUICK CONTROL

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

The file-server evidence checklist

- ☐ Day-one preservation email sent: snapshots, logs and backups held; confirmation received and filed
- ☐ System biography documented: platform, migrations, restores, sync, admin model
- ☐ Audit configuration and permission structure exported early, as exhibits
- ☐ Scope set by named paths, dates and bracketing snapshot states, with the broad-preservation backstop
- ☐ Acquisition by snapshot or protocol-run live collection; no admin drag-and-drop
- ☐ Content hashed and manifested; counts reconciled; samples verified
- ☐ History layer collected: snapshot extracts, differencing, backup generations as needed
- ☐ Attribution planned as four links; implicated endpoints collected alongside the server
- ☐ Service/admin activity separated from human activity in every chronology
- ☐ All three layers delivered under one custody record, with limitations stated

Red flags

- "IT copied the folder across for you."
- Snapshot pruning still running weeks into the dispute.
- A schedule of "modifications" clustering on a migration or restore date.
- Attribution built on the modified-by column alone.
- An opponent silent on audit configuration while asserting who accessed what.
- Permission changes on disputed trees after proceedings began.
- Findings citing snapshots nobody preserved.

When to involve a digital forensic expert

On day one, to shape the preservation email and establish the estate's realities (platform, schedules, audit posture) before they expire; at scoping, where the snapshot-versus-live-versus-image decision and the three-layer plan are the examiner's product; at collection, per the playbook; and at analysis, where server literacy earns its fee: differencing chronologies, migration-aware date reading, and four-link attribution reported under CPR Part 35. Against an opponent, the same literacy runs in reverse: their server-based schedule of sins is tested against the biography, the audit ceiling and the service-account noise, and: as Examples 2 and 3 show: shared-drive accusations built without this guide's disciplines tend not to survive contact with them.

§ 12 · COMMON QUESTIONS

Frequently asked questions

Can we just have IT export the relevant folders to save cost?

IT can assist; the export must still be forensic: metadata-preserving tooling, hashing, manifests, logging, under an examiner-written method with supervision and sampling (the hybrid conditions this series keeps insisting on). The drag-and-drop alternative rewrites timestamps across the very documents in dispute and converts a defensible collection into a witness-statement apology. The cost difference between the two is small; the evidential difference is the whole case's metadata.

The files we need were deleted from the share months ago. Are they gone?

Ask the history layer, quickly: snapshots reach days to weeks; previous-versions data similar; backup generations months to years, until recycling eats the relevant generation. The §10 suspension email is therefore the first act, today; the second is establishing which states bracket the deletion; recovery then usually follows. What rarely helps on servers is carving: enterprise storage abstractions bury released space: so the schedule conversation is the recovery conversation.

Can the server tell us who read a document, not just who changed it?

Only if read-auditing was configured for that tree during the period, which is the exception rather than the rule (read events are voluminous and often left off). Where it was on, the log answers directly; where it was off, "who read it" is rebuilt from the reader's side: endpoint recency artefacts, application traces, and (for remote access) VPN and authentication correlation: per guide 41. The audit-configuration export tells you which world you are in before anyone promises the client an answer.

Our opponent says their server "only keeps thirty days of logs" so nothing can be shown. Is that the end of it?

It is the beginning of three questions: what did the thirty-day window overlap (rotation schedules preserved from your first preservation letter onward: late suspension is itself a disclosure issue); what do snapshots and backups hold (log files are files: earlier generations often contain earlier log states); and what do the endpoints and adjacent systems record (authentication servers, VPNs, the implicated workstations). Short retention narrows one source; it does not answer for the estate, and a Part 35 review of what should exist elsewhere usually reframes the correspondence.

The share lives in SharePoint/OneDrive now. Does this guide still apply?

The three-layer thinking applies exactly; the machinery differs: cloud platforms keep their own version histories, recycle stages, audit logs and permission models, generally richer and longer-retained than on-premises defaults, collected through platform mechanisms rather than snapshots. The dedicated Microsoft 365 and cloud guides in this series carry the specifics; the transferable instruction is unchanged: name content, history and context in every preservation notice, whichever world hosts the share.

How disruptive is server collection to the business?

With snapshots, negligible: the freeze is instant and collection proceeds from the frozen copy while production runs on. Live share collection is scheduled into quiet windows and throttled; full imaging, where genuinely needed, is negotiated like any guide 39 disruption. The honest cost is administrative attention (IT time for snapshots, exports and questions) rather than downtime: which is one more reason the cooperative, planned collection beats the improvised one on every axis this series measures.

§ 13 · REFERENCE

Glossary

Audit configuration
Which file events the system was set to log, where, since when; the ceiling on actuality evidence.

Backup generation
A dated backup state; the long-reach history layer, eaten by recycling schedules.

Differencing
Comparing successive snapshot states to derive a tree's change chronology.

Four-link attribution
File event → account → session/endpoint → person; every link evidenced separately.

Permission structure
The folder-level access map: who could reach what; the map of possibility.

Previous versions
Snapshot-backed earlier file states exposed by Windows servers.

Service account
Non-human account (backup, sync, migration) whose file activity must be separated from users'.

Share-level collection
Metadata-preserving forensic acquisition of defined paths from a live system.

Snapshot
A point-in-time frozen state of a volume or VM; the acquisition moment of choice.

System biography
The share's history of migrations, restores and reorganisations; the date-literacy layer.

Sources and authoritative references

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

- cflab.uk, digital forensics services (server and live-system acquisition, snapshot examination, attribution analysis, CPR Part 35 reporting): cflab.uk/digital-forensics-services · guides library: cflab.uk/guides
- e-discovery.uk, EDRM Phase 02 · Preservation and Phase 03 · Collection (three-layer server collection at disclosure scale; "defensibility is built, not retrofitted"): e-discovery.uk/edrm/preservation · e-discovery.uk/edrm/collection
- e-discovery.uk, practice method and Knowledge Centre: e-discovery.uk · e-discovery.uk/knowledge
- NPCC / College of Policing digital evidence guidance: college.police.uk, digital devices guidance · ISO/IEC 27037:2012: iso.org, 27037
- CPR Part 31 / PD 31B and PD 57AD (preservation and the DRD): justice.gov.uk, Part 31 · justice.gov.uk, PD 57AD
- CPR Part 35: justice.gov.uk, Part 35 · ICO, UK GDPR guidance: ico.org.uk

§ HOW A SPECIALIST LABORATORY CAN ASSIST

Working with Computer Forensics Lab

Server collection is standing laboratory work: the day-one preservation call with your client's IT, the estate reconnaissance, snapshot-based and live-share acquisition to protocol, and the three-layer deliverable: content manifested and hashed, history differenced, context exported: under one custody record. Analysis follows with server literacy: migration-aware date reading, service-account separation, four-link attribution with endpoint correlation, and CPR Part 35 reporting that states audit ceilings and snapshot intervals as plainly as findings. Against opposing server claims, we run the same disciplines in reverse. Through **e-discovery.uk**, server collections flow into processing and review beside every other source, scoped so you pay for substance, not volume. Conflicts checks, confidential scoping discussions and fixed-fee estimates are routine; the preservation email can be settled with an examiner today.

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