COMPUTER FORENSICS LAB

London digital forensics & e-discovery · Established 2007

Domain, DNS and Hosting Records

Who Is Behind a Website, and How Its Infrastructure History Is Proved

Fraudulent shops, copycat brands, phishing domains, anonymous defamation sites, spoofed invoices from lookalike addresses: the first question in each is who is behind this domain, and the second is what its infrastructure looked like at the material time. The answers live in registration records constrained by post-GDPR redaction, DNS configurations and their archived histories, hosting and registrar accounts with payment trails, certificate transparency logs that publish every TLS certificate ever issued, and email authentication records that decide spoofing disputes. This guide maps the infrastructure evidence layer for UK lawyers: what each record proves, the unmasking sequence from redacted WHOIS to a named registrant, historical reconstruction of where a site pointed and when, the registrar and host disclosure routes, and the join to the site-content and web-log evidence of the neighbouring guides.

⏱ Estimated reading time: 25 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. Infrastructure investigation runs through its fraud, brand-protection and defamation casework: registrant unmasking, DNS and hosting history reconstruction, certificate-transparency analysis, spoofed-email authentication work and evidence packages supporting registrar, host and Nominet processes. Its eDiscovery arm, **e-discovery.uk**, hosts infrastructure chronologies beside the documentary estate.

ESTABLISHED 2007 · LONDON | ISO 17025-ALIGNED PROCEDURES | DOMAIN & DNS INVESTIGATION | REGISTRANT UNMASKING | CPR PART 35 EXPERT REPORTS | FULL CHAIN-OF-CUSTODY DOCUMENTATION

§ CONTENTS

In this guide

01 Executive summary

02 The problem in plain English: the address book of the internet

03 The record layers: registration, DNS, hosting and certificates

04 Unmasking: from redacted WHOIS to a named registrant

05 Reconstructing history: what the domain did and when

06 Email authentication: SPF, DKIM, DMARC and spoofing disputes

07 Remedies and process: registrars, hosts, Nominet and the courts

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: PUBLIC INFRASTRUCTURE RECORDS NARROW THE FIELD; REGISTRAR AND HOST DISCLOSURE NAMES THE HAND: AND THE HISTORY IS ARCHIVED EVEN WHEN THE SITE IS GONE

Domain evidence works in two movements. **Narrowing:** the public layers: registration data (largely redacted since GDPR, but with registrar, dates and status intact), current and archived DNS, hosting and IP neighbourhoods, certificate transparency logs and the site's own technical fingerprints (analytics IDs, verification codes) reused across an operator's estate: often shrink "anonymous" to one identifiable operator running many domains, all documented from records anyone can verify. **Naming:** the private layers: registrar account details, host subscriber records, payment instruments: reached by disclosure request, abuse process, Norwich Pharmacal order or Nominet/UDRP procedure, converting the narrowed field into a person. Meanwhile the archive layer (historical WHOIS and DNS services, certificate logs, the Wayback Machine per guide 79 §5) reconstructs what the domain pointed at and served at the material time, so takedown does not equal escape. Everything captures with dates, sources and hashes: infrastructure evidence is public, but its evidential form is made, not found.

- **Redaction hides, it does not erase:** post-GDPR WHOIS still yields registrar, dates, status and nameservers: and the registrar holds the rest for the court.
- **Fingerprints cluster domains:** shared analytics IDs, certificates, nameservers and hosting group a fraud estate under one operator.
- **History is archived by third parties:** passive DNS, WHOIS history and certificate logs date what the domain did: takedowns arrive too late to help the wrongdoer.
- **Email spoofing is decided by authentication records:** SPF/DKIM/DMARC results in headers separate forged mail from compromised accounts.
- **Capture evidentially:** dated, sourced, hashed lookups: the public record turned into an exhibit properly.

§ 0 2 · FIRST PRINCIPLES

The problem in plain English: the address book of the internet

A domain name is an entry in a public address book: registered through a registrar, pointed by DNS records at hosting infrastructure, and dressed with certificates so browsers trust it. None of that can be done invisibly: registration creates records, pointing creates DNS entries that third parties continuously archive, certificates are published by design in transparency logs, and hosting means a commercial relationship with a provider who invoices someone. The wrongdoer's site may vanish overnight; the paper trail of its existence: who registered, where it pointed, what it served, who paid: substantially cannot.

What changed in 2018 is the front door: public WHOIS once named registrants; GDPR redacted most of it, and "who owns this domain" now takes a sequence rather than a lookup. The sequence is this guide: harvest what remains public, cluster the operator's estate by its reused fingerprints, reconstruct the history from the archives, and take the narrowed, documented package to the registrar, the host or the court for the name. It works routinely: redaction moved the name behind a disclosure step; it did not delete it.

§ 0 3 · THE LAYERS

The record layers: registration, DNS, hosting and certificates

- **Registration (WHOIS/RDAP):** registrar, creation, update and expiry dates, status codes and nameservers survive redaction: creation date alone decides many cases (the "established supplier" registered eleven days before the invoice), and RDAP is the structured modern lookup.
- **DNS records:** A/AAAA (where the site lives), MX (who handles its mail), NS (who controls its zone), TXT (verification codes and email policies): the current configuration read as a statement of the operator's arrangements.
- **Hosting and IP context:** which provider, which ASN, what else sits on the address or its neighbourhood: shared infrastructure suggesting shared operation, weighted honestly against the reality of mass shared hosting.
- **Certificate transparency:** every TLS certificate publicly logged with domains, issuance dates and issuers: revealing subdomains, sibling domains sharing certificates, and preparation timelines (the certificate issued three weeks before the phishing wave dates the plan).
- **Site-level fingerprints:** analytics and tag IDs, ad-network codes, verification TXT records, favicon and template hashes: the reuse layer that clusters one operator's many domains, per §4's method.

§ 0 4 · THE SEQUENCE

Unmasking: from redacted WHOIS to a named registrant

- **Harvest the public residue:** registrar, dates, nameservers, status; any unredacted fields; the RDAP abuse contacts that route later steps.
- **Cluster the estate:** pivot on fingerprints: shared certificates, analytics IDs, nameserver patterns, registration timing, template reuse: one domain becomes an operator's mapped portfolio, and portfolios leak (one domain in forty registered before redaction, or with a real address in a payment page, names the lot).
- **Content-side identifiers:** the sites themselves: contact details, company numbers, payment recipients, social handles: cross-checked against Companies House and the correspondence estate: wrongdoers economise on identities.
- **Provider disclosure:** registrar and host abuse and legal channels, UK GDPR-aware requests, and where cooperation fails, Norwich Pharmacal orders against the registrar, host or payment processor: drafted with the precision guide 79 §6 demands of ISP requests.
- **The rung stated:** "the same operator runs these nineteen domains" is a clustering finding; "that operator is X" needs the disclosure or content rung: pleaded per the series' standing attribution rule.

§ 0 5 · THE PAST

Reconstructing history: what the domain did and when

- **WHOIS history services:** commercial archives holding pre-redaction and interim records: ownership changes, contact details from older eras, the transfer dated to the week the dispute began.
- **Passive DNS:** resolution archives recording which IPs a name pointed to over time: the "it was pointing at their server in March" answer, source and observation dates cited.
- **Certificate timelines:** issuance and renewal history dating when the domain was dressed for use, and when siblings joined the estate.
- **Content archives:** the Wayback Machine and monitoring services for what the site said, joined per guide 79 §5's page-state chorus: infrastructure history for where it lived, content history for what it showed.
- **Registrar and host records:** account logs, transfer histories, payment records: the authoritative past, reached by the §4 disclosure routes when the archives need confirming.

§ 0 6 · THE MAIL QUESTION

Email authentication: SPF, DKIM, DMARC and spoofing disputes

- **The framework in one breath:** SPF lists the servers allowed to send for a domain; DKIM signs messages cryptographically; DMARC tells recipients what to do with failures and where to report: all published as DNS records, all checkable at any date via DNS history.
- **Reading the headers:** a received message's Authentication-Results record the checks at delivery: pass/fail per mechanism: the first exhibit in every spoofed-invoice case, read with guide 53's header disciplines.
- **Spoof vs compromise:** the fraud-payment fork: authentication failure plus foreign infrastructure suggests external spoofing or a lookalike domain; clean passes from the genuine domain point at account compromise (guide 79 §6's takeover analysis) or an insider: the fork decides who gets sued.
- **Lookalike analysis:** the cousin domain (hyphens, transposed letters, alternative TLDs) with its own §3 record set: creation date, certificate, MX arrangements: registered days before the fraud and evidencing the plan.
- **The policy posture as negligence evidence:** a payer or payee's absent/weak DMARC posture at the material time (provable from DNS history) increasingly features in loss-allocation arguments after push-payment frauds: checked for your own client early.

§ 0 7 · GETTING IT STOPPED, GETTING IT NAMED

Remedies and process: registrars, hosts, Nominet and the courts

- **Abuse channels:** registrar and host abuse processes for fraud and phishing: fast, evidence-hungry, and fed by the documented package this guide builds: takedown first, litigation preserved by prior capture.
- **Nominet DRS and UDRP:** the domain-dispute procedures for .uk and gTLDs: rights, confusing similarity, bad faith: with the §3-§5 evidence (registration timing, use history, estate clustering) as the factual core.
- **Court routes:** Norwich Pharmacal against registrars, hosts and processors for identity; injunctions including transfer and blocking orders; the infrastructure chronology as the supporting evidence throughout.
- **Preservation before takedown:** the standing sequence: capture the site, the DNS, the certificates and the archives evidentially before the abuse report succeeds: your own remedy must not destroy your own evidence.
- **Cross-border reality:** registrars, hosts and registrants scatter across jurisdictions: routes chosen for where the levers are (the .uk registry, the UK-reachable processor), with expectations set honestly.

§ 0 8 · THE WIDER MAP

Source architecture: where else the evidence lives

Per this series' source-architecture discipline: the domain's public records are one stratum of a wider estate. The evidence also lives in: the registrar's account records (identity, logins, payment methods, transfer history); the host's subscriber and server records (and the site files and logs themselves, guide 79's territory); payment processors behind the site's checkout and the registrar's invoices; the third-party archives (WHOIS history, passive DNS, certificate logs, content archives) that outlive takedowns; the victim-side record (the phishing email with its headers, the browser history of the visit, the bank's records of where money went); connected platform accounts (the social profiles, marketplace listings and ad accounts the domain advertised through, each with its own operator records); and the correspondence estate around the scheme. When the domain is abandoned and its records purged, the archives date its life, the processors follow its money, and the platform accounts name its operator.

EVIDENCE	PUBLIC RECORDS (LIVE)	THIRD-PARTY ARCHIVES	REGISTRAR / HOST ACCOUNTS	PAYMENT LAYER	VICTIM-SIDE RECORDS	DELETED / RECOVERABLE
Who registered	Redacted residue	WHOIS history	Y: the name, by disclosure	Card/PayPal identity	n/a	Archives + disclosure rebuild
Where it pointed	Current DNS	Passive DNS: the timeline	Zone change logs	n/a	Visited-IP artefacts	Y: archived by design
What it served	The live site	Content archives	Server files + logs	Checkout records	Screenshots, cache, emails	Guide 79 §5's chorus
When it was built	Creation date, cert dates	Cert transparency timeline	Account history	First transactions	First contact dates	Y: published logs persist
The estate's extent	Fingerprint clustering	Cross-archive pivots	Same-account portfolios	Same-instrument links	Parallel victim reports	Clusters survive takedowns

Fallback strategy in one line: when the domain dies, work the archives for its life, the processors for its money and the platform accounts for its operator: infrastructure is disposable; its records are not.

§ 0 9 · IN THE WILD

Worked examples

EXAMPLE 1 · THE SUPPLIER ELEVEN DAYS OLD

A manufacturer paid £214,000 against an email chain from its long-standing German supplier: the same names, the same signature blocks, a "new account" for the payment. The infrastructure record unravelled it in an afternoon: the sending domain was a cousin (one transposed letter), created eleven days before the first email; its certificate was issued the same week; its MX records pointed at a free mail platform the real supplier had never used; and the genuine supplier's DMARC policy: which would have flagged the fraud: turned out to be absent, a fact now central to the loss-allocation argument with the bank. Passive DNS showed the cousin domain parked at the same host as four other cousins of unrelated companies: the §4 clustering that turned one fraud into an operator profile for the NPCC referral.

EXAMPLE 2 · THE "COMPETITOR" REVIEW SITE

A chain of veterinary clinics was savaged by an anonymous "independent review" site praising one rival in every town. Public records built the attribution ladder: the site's pages carried a Google Analytics ID shared with the rival's own marketing site; certificate transparency showed both domains on one multi-domain certificate for a four-month period; historical DNS placed them on the same small agency's server; and the review site's domain had been registered three days after the rival's franchise agreement (disclosed later) was signed. The Norwich Pharmacal order against the registrar returned the rival's marketing director as registrant, but by then the clustering evidence had done the persuading: the defamation claim settled with the site's transfer and a contribution to costs, the shared analytics ID quoted in the settlement recitals.

EXAMPLE 3 · THE DOMAIN THAT CHANGED HANDS MID-DISPUTE

A distributor terminated for cause continued trading from the brand's lookalike domain; by the time proceedings issued, the domain's records showed a new "unrelated" registrant in another jurisdiction and the defence pleaded coincidence. The archives disagreed: WHOIS history dated the transfer to nine days after the letter before action; passive DNS showed the site never moved servers through the "change of ownership"; the certificate was reissued to the same organisation identifier; and the site's checkout continued paying the same processor account, confirmed by the processor's Norwich Pharmacal response. The transfer was a costume change, dated and documented. The injunction extended to the transferee, and the transfer's timing anchored the bad-faith finding in the parallel Nominet DRS.

§ 10 · WHERE IT GOES WRONG

Common mistakes and technical limitations

Common mistakes

- **Stopping at redacted WHOIS:** "the registrant is hidden" pleaded as a dead end: §4's sequence never run.
- **Takedown before capture:** the abuse report succeeding against your own evidence: the standing sequencing error.
- **Screenshots without provenance:** undated, unsourced lookups exhibited: captured properly (tool, date, hash) or challenged easily.
- **Shared hosting over-read:** same-IP inferences on mass shared infrastructure: thousands of innocent neighbours ignored.
- **Header analysis skipped:** spoofing pleaded without Authentication-Results: the spoof/compromise fork unexamined, the wrong defendant sued.
- **The client's own posture unchecked:** the absent DMARC record discovered by the opponent first: Example 1's second front.
- **Disclosure requests drafted loosely:** registrar and processor requests without the precision (domains, dates, account identifiers) that gets them answered.

Technical limitations

- **Privacy services and false details:** registrant data can be proxied or fabricated: the payment layer usually resists fabrication better than the contact form.
- **Archive coverage is sampled:** passive DNS and content archives record observations, not continuous truth: findings cite observation dates and gaps.
- **Fingerprints can mislead:** agencies reuse analytics IDs across unrelated clients; templates are bought by strangers: clustering findings weigh alternatives per the standing fairness rule.
- **Jurisdiction limits reach:** some registrars and hosts answer no civil process usefully: routes chosen for reachable levers, expectations set early.
- **DNS answers vary by vantage:** geo-split and fast-flux configurations serve different answers to different askers: captures note the resolver and vantage used.

§ 11 · INTERROGATORIES & DRAFTING AIDS

Questions to ask · Suggested wording

Ask your client

- What exactly arrived: the emails with full headers, the URLs visited, the payment instructions: preserved per guide 53 and guide 79's counterpart disciplines?
- What is our own domain posture: SPF, DKIM, DMARC: now and at the material time?
- Which remedies matter most: money, takedown, identity, transfer: so the §7 routes are sequenced to serve them?

Ask your opponent

- Please identify all domains registered, controlled or used by you or on your behalf relating to [the issues], including through privacy or proxy services, with registration dates, registrars and hosting arrangements.
- Please produce your registrar and hosting account records for [domains], including account identities, payment methods, DNS change logs and transfer records for [period].
- Please state whether [domain] or its content has been transferred, modified or taken down since [date], and by whom, producing the records of any such steps.

Ask your eDiscovery / forensic provider

- What does the public-and-archive package establish, with what dates and sources: and what does it not?
- Which disclosure target (registrar, host, processor, platform) is the fastest route to a name here?
- What did the authentication records say at delivery: and which side of the spoof/compromise fork are we on?

SUGGESTED WORDING · INFRASTRUCTURE LIMB FOR THE PRESERVATION LETTER

> "Your client must immediately preserve all records relating to the registration, control and operation of [domains], including: (a) registrar and hosting account records, invoices, payment records and correspondence; (b) DNS zone files and change logs; (c) website files, databases and server logs (as to which see the web-log limb of this letter); (d) certificates and keys; and (e) records of any transfer, sale or delegation of the domains. No domain shall be transferred, allowed to lapse, or have its records deleted pending determination of these issues. We put you on notice that historical registration, DNS and certificate records for these domains have been independently captured and preserved."

§ 12 · QUICK CONTROL

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

The infrastructure checklist

- ☐ Public residue harvested: registrar, dates, status, nameservers, RDAP contacts
- ☐ Everything captured evidentially: tool, date, vantage, hash
- ☐ Estate clustered on fingerprints; alternatives weighed
- ☐ Histories pulled: WHOIS archives, passive DNS, certificate logs, content archives
- ☐ Headers and Authentication-Results analysed in any mail-borne matter
- ☐ The client's own SPF/DKIM/DMARC posture documented, then fixed
- ☐ Capture completed before takedown or DRS steps
- ☐ Disclosure requests drafted to registrar/host/processor precision
- ☐ Cross-border strategy set by reachable levers
- ☐ Findings pleaded to their rung: clustered, disclosed, or named

Red flags

- Domains created days before the transactions they carried: Example 1's eleven.
- Cousin spellings and alternative TLDs of trusted counterparties.
- Transfers, registrant changes and takedowns dated to the letter before action: Example 3.
- Shared analytics IDs and certificates behind "unrelated" sites: Example 2.
- MX records routing a corporate brand through anonymous mail services.
- Authentication failures in the headers of payment-changing emails.
- Privacy-proxied estates with synchronised registration bursts.

When to involve a digital forensic expert

At first sight of a suspect domain, where evidential capture and the archive pulls must precede takedown; at analysis, where clustering, history reconstruction and header authentication build the package (Examples 1-3's machinery); at the procedural steps, where registrar, host and processor disclosure requests and Nominet/UDRP evidence need precision; and at reporting under CPR Part 35 with sources, observation dates and rungs stated. Through **e-discovery.uk**, infrastructure chronologies sit in review beside the correspondence and the web-log record per guides 70 and 79.

§ 1 3 · COMMON QUESTIONS

Frequently asked questions

WHOIS shows the registrant as "REDACTED FOR PRIVACY". Dead end?

Starting line: the residue (registrar, dates, nameservers) plus fingerprint clustering, archives and content identifiers usually narrows sharply, and the registrar holds the actual identity for disclosure by request or order: §4's sequence exists precisely because the lookup alone stopped working in 2018. Treat redaction as one extra procedural step, not an outcome.

The fraudulent site has already been taken down. Have we lost the evidence?

Mostly no: registration and DNS history, certificate logs and content archives were recording all along, the registrar and host retain account records, and the payment layer retains the money's path: §8's map is the recovery plan. What takedown does cost is anything never archived: which is why on live matters capture always precedes the abuse report.

Can we find every domain a wrongdoer operates?

Often most of them: certificate transparency, shared fingerprints, nameserver and registration patterns, and same-account disclosure from registrars map estates well: Example 1's operator ran to five cousins, Example 2's to two. "Every" is unprovable: findings say "the following linked domains", with the linking evidence per domain, and stay open to additions.

An email claims to be from our client's domain. How do we prove it was not?

From the receipt evidence: the message's Authentication-Results and full headers show whether SPF/DKIM/DMARC passed at delivery and which infrastructure actually sent it; the client's published records at the date (from DNS history) show what should have passed; and the sending domain's own §3 profile (cousin spelling, creation date, MX) usually completes the story. Where everything passed cleanly from the genuine domain, the inquiry pivots honestly to compromise or insider: guide 79 §6's fork.

Is a Norwich Pharmacal order really available against a registrar or payment processor?

In principle yes, on the standard mixed-up-in-wrongdoing basis, and processors are often the most productive target because payment instruments resist fabrication: the practical constraints are jurisdiction and drafting precision (exact domains, accounts, dates, in the respondent's own vocabulary). The §4 package is what persuades the court the order will actually identify someone.

Should our client publish DMARC even if it is the victim here?

Immediately: a reject-policy DMARC record (rolled out carefully) makes the client's domain hard to spoof and generates reports of attempts: prospectively protective, and retrospectively relevant, since the absent policy at the material time is increasingly raised in loss-allocation arguments after payment frauds, as Example 1's bank correspondence shows. It is the rare remediation that is also good evidence hygiene.

§ 14 · REFERENCE

Glossary

ASN
Autonomous system number; the network operator behind an IP range.

Certificate transparency
The public logs recording every issued TLS certificate.

DMARC
The published policy telling recipients how to treat authentication failures.

DKIM
Cryptographic signing of outbound mail, verifiable by recipients.

Nominet DRS
The dispute resolution service for .uk domains.

Passive DNS
Archives of observed DNS resolutions over time.

RDAP
The structured successor to WHOIS lookups.

Registrar
The company through which a domain is registered; holder of the identity.

SPF
The DNS record listing servers authorised to send a domain's mail.

UDRP
The uniform dispute policy for gTLD domain disputes.

Sources and authoritative references

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

- cflab.uk, digital forensics services (domain investigation, unmasking, spoofing analysis, CPR Part 35 reporting): cflab.uk/digital-forensics-services · guides library: cflab.uk/guides
- e-discovery.uk, EDRM Phase 01 · Identification, Phase 02 · Preservation and Phase 03 · Collection (infrastructure sources in the data map; capture-before-takedown; evidential acquisition): e-discovery.uk/edrm/identification · 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
- Nominet, .uk Dispute Resolution Service: nominet.uk, DRS · ICANN, UDRP: icann.org, UDRP
- NCSC, email security and anti-spoofing guidance (SPF, DKIM, DMARC): ncsc.gov.uk, email security
- PD 57AD, CPR Part 31 and CPR Part 35: justice.gov.uk, PD 57AD · justice.gov.uk, Part 35 · ICO, UK GDPR guidance: ico.org.uk

§ HOW A SPECIALIST LABORATORY CAN ASSIST

Working with Computer Forensics Lab

Infrastructure instructions at the laboratory run the two movements of §1: the public-and-archive package first: captured evidentially, clustered on fingerprints, historied through WHOIS, passive DNS and certificate logs: then the naming step through the registrar, host, processor or court, with requests drafted to the precision that gets them answered. Mail-borne frauds get the header and authentication analysis of §6, including the client's own posture and its remediation; takedowns are sequenced after capture, always; and everything reports under CPR Part 35 with sources, dates and rungs stated, per Examples 1-3. Through **e-discovery.uk**, the infrastructure chronology sits beside the correspondence and web-log record on one reviewable timeline. Conflicts checks, confidential scoping discussions and fixed-fee estimates are routine; and if a suspect domain is live right now, the evidential capture happens before anyone reports it: today, in that order.

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