What a tamper-proof log still cannot tell you

Agent traceability work has produced records that cannot be altered. It has not produced records that can say which law was in play, because the signal that would answer that is the one every privacy instinct says not to keep. The instinct is right about the raw signal and wrong about the field.

1 The gap in agent traceability records

Work on agent traceability has converged fast on a record of what an agent did: one entry per action, refusals included, hash chained, signed, and registered somewhere the operator cannot quietly revise. That is a real achievement and it is close to finished. Read the record shapes now in circulation against the things law can require a record to hold, and integrity is the best-served category by a distance.

The same reading finds one category served by almost nothing: which person was affected, and which place they were in. Of the shapes surveyed, one carries a single country code with no role attached and no evidence behind it, one excludes end-user identity outright and correctly for its own purposes, and none carries the law that was in play. The gap is not an oversight in any one document. It falls between two layers: the record format treats a place as identity and keeps it out, and the governance framework above treats the record format as having handled it.

> THE SHAPE OF THE PROBLEM
> A tamper-evident record of an action nobody can attribute to a legal regime proves that something happened and not that it was allowed.

2 How jurisdiction determines applicable law

The reason the field matters is not that regulators want to know where anybody lives. It is that the trigger in almost every instrument is conduct directed at a place. Article 3(2) of the GDPR reaches a controller with no establishment in the Union where its processing relates to offering goods or services to people in the Union, or to monitoring their behaviour there; Recital 23 puts the test in what the controller shows about its own intent rather than in where a website happens to be reachable from (Regulation (EU) 2016/679). Article 2(1)(c) of the EU's AI law reaches a provider or deployer in a third country where the output produced by the system is used in the Union (Regulation (EU) 2024/1689). The DSA, the UK's online safety regime and the US state age-verification statutes each run on material being made available to people in a place. American personal-jurisdiction doctrine, from *Zippo* through *Ford Motor Co. v. Montana*, asks a version of the same question about availing yourself of a market.

So a record that shows an agent acted correctly under a rule has to be able to say which rule was in force for that action, and that is a question about a place. A log that cannot answer it can show that a policy was applied. It cannot show that it was the right policy.

3 Location data, personal data, and what to record

Two different things are called an address in this area, so this paper never writes the word bare: an **IP address** is the network identifier a request arrives with, and a **residential address** is what somebody enters into an account. They answer different questions and only one of them is a coordinate.

The reason the field is usually left empty is a sound one. A client IP address is personal data wherever the party holding it has the legal means to link it to a person, which is the holding in *Breyer* (C-582/14), and the GDPR's own Recital 26 asks whether a value allows singling out by anyone reasonably likely to try. California's privacy statute makes geolocation sensitive once it places a person inside a circle of 1,850 feet. UK rules on electronic communications gate network-derived location data more tightly still. A record built to outlive the incident it describes is the wrong place for any of that.

A digest does not rescue it, and this is the part most often got wrong. A record carrying a hash of an IP address carries that address: a space of about 4,300,000,000 values is enumerated in seconds. The IETF draft for agent action records states the rule for its own equivalent problem in as many words: hashing is not anonymisation for low-entropy identifiers, and identity must be excluded rather than digested (draft-mih-scitt-agent-action-capsule-04, section 14.1).

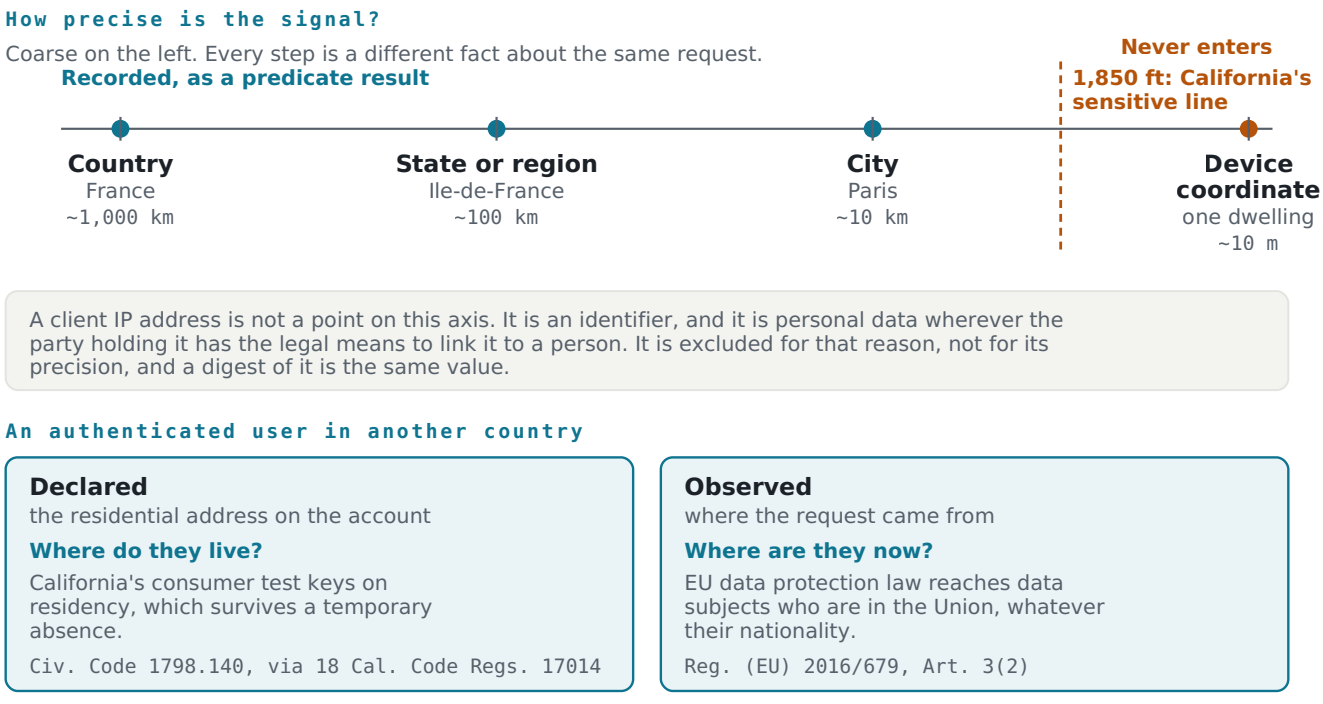

Figure 1. The distances are orders of magnitude, not a claim that any one lookup resolves to a stated radius; what the axis asserts is the order of the steps and which side of California's line each falls on. The client IP address sits off the axis deliberately, because excluding it for being too precise teaches the wrong rule: it is excluded for being an identifier, which is also why a digest of it is the same value.

What does not follow is the conclusion usually drawn from it. "The IP address is personal data, so the record carries nothing about place" discards the answer along with the evidence. The two are separable, and separating them is the design:

Tier	What it holds	Example
Clear	A derived, non-identifying result: the output of a deterministic check, at country or state granularity, with *unknown* a permitted value.	`directed_at: urn:ungovr:eu`
Commitment	A binding to an evidence record held elsewhere, which includes a random value of its own so the commitment is inert to anyone not shown the evidence.	`evidence: <commitment>`
Never enters	An IP address, a precise coordinate, an end-user or session identifier, in clear or as a bare digest.	nothing

The clear field needs a vocabulary, or every operator invents one and no verifier can read two records side by side. We propose the identifier LexLint already publishes for a jurisdiction: `urn:ungovr:<slug>`, so the European Union is `urn:ungovr:eu`, France is `urn:ungovr:fr`, California is `urn:ungovr:us/ca` and the City of San Francisco is `urn:ungovr:us/ca/san-francisco`. It nests, which matters because the answer is often a country and sometimes a city; it is resolvable, so a reader can look up what was claimed; and it is published under Creative Commons Attribution 4.0, so a standards body or a competitor can adopt it without asking us. An identifier nobody else may use is not an interoperable field.

A country or state label sits far outside California's 1,850-foot line by construction, and it is the granularity regulators already accept as compliance tooling: US sanctions enforcement treats coarse location screening as an expected control rather than as surveillance, and the state age-verification statutes require exactly this shape, locate and characterise and then discard, as a matter of law rather than as good practice.

One caveat belongs beside the table, because a field is safe only in the company it keeps. A country label read alone is a derived, non-identifying value. A country label sitting next to a session identifier in the same record is a country label attached to a person. The rule is enforced across every field admitted to the record, not checked once on this one. The statutory reading this table rests on, jurisdiction by jurisdiction, is Knowing which law applied, without making the log personal data.

4 Identifying the real client IP address

There is a failure between the two halves above that neither half catches, and it is the reason a design that is correct on paper produces a record that is wrong in production.

At an origin behind a content delivery network, a load balancer or any reverse proxy, the peer IP address on the socket is the edge node's, not the user's. A jurisdiction check run on it resolves the point of presence that relayed

the request. It returns a well-formed country code, on time, with no error. It is simply the wrong country, and nothing about the value says so. A user in Bavaria, reaching an origin in Virginia through an edge node in Frankfurt, yields three countries, and the only one that answers any legal question is the one the socket does not carry.

Server IP addresses carry no jurisdiction signal at all. Where an origin sits is an infrastructure decision, and the establishment questions that do turn on an operator's own footprint are answered from corporate records rather than from a route. The client's IP address is the only one in the exchange that bears on where a duty came from, and behind an edge it reaches the origin in one way only: as an assertion the edge makes in a forwarded header, over a connection the origin can authenticate as its own. That makes the trust boundary part of the evidence rather than a deployment detail, and an evidence record that does not say which signal the check ran on, and on what basis the origin believed it, is not checkable later.

One request, three hops

	1. User to edge recorded in the CDN's log	2. Edge to origin recorded in the origin's log	3. Origin to agent service recorded in the agent service's log
Source IP	the user's own	the edge node	the origin
Destination IP	the edge node	the origin	the agent service
What that log can see of the user's place	Directly: the user's IP address is on the wire	Only what the forwarded header asserts	Nothing: it must be told
What that log records instead	fr, derived here, and the IP address itself dropped	fr, as forwarded, plus the origin's own establishment	fr, as it was told, plus its own establishment

The user's own IP address appears at exactly one hop, and that is the hop whose log must not keep it. Everything downstream is working from what it was told, so the jurisdiction is determined once, where the signal exists, and travels with the request. A declared residential address, where an account carries one, reaches no hop at all and has to be passed along the same way.

Figure 2. Drawn as hops rather than as parties, because "the log" is not one log: each hop is written down by whoever terminates it, and the end user writes none at all. The third hop is a third-party service, which has an establishment of its own and no sight of the user's IP address in any form. A second machine belonging to the operator would share the operator's establishment and still see nothing. The bottom row is the answer those IP addresses were being asked for: every log carries the same jurisdiction for the user, and the two server-side logs carry their own establishment beside it.

That has a consequence the record alone does not cover: **the reading has to travel**. The edge is the only place the real client IP address exists, so every party downstream of it is working from what it is told. An origin can read the forwarded header; a third-party agent service called by that origin cannot, because the signal never reaches it. If the answer is not passed along, a service acting on a French user's request applies whatever regime it would have applied to anyone, and logs it that way.

So the jurisdiction result belongs in the call, not only in the log. A processor acting on a controller's documented instructions needs those instructions to say which law is in play, because it has no way to work it out; and each layer then records the reading it was given rather than deriving its own, which is also what makes two records of one action agree. The evidence commitment travels with it, so a layer that is handed an answer can still show what the answer was checked against.

> WHY THIS IS WORSE THAN AN EMPTY FIELD
>
> A wrong jurisdiction label committed to an append-only, signed, hash chained record is worse than an absent one. Every integrity property the format was built for still holds, and holds around a false answer. The record verifies. The chain is intact. The signature checks. Nothing downstream has any reason to doubt a field that passes every test the format knows how to run, which is why tamper evidence makes this failure harder to find rather than easier.

Two rules follow, and they are cheap. The predicate names the signal it ran on, so a later reader can tell a reading from an artefact of the network path. And *unknown* stays a permitted result: a check that cannot answer says so, because coercing an unresolved signal to a default country is how a wrong answer is manufactured rather than merely inherited.

5 Jurisdiction and incident reporting deadlines

The jurisdiction question is usually filed under access control, which undersells it. Jurisdiction decides which regulator is owed a report and how long you have. Article 33 of the GDPR runs 72 hours from awareness. The NIS2 Directive, the CRA, Article 73 of the EU AI Act and the SEC's Form 8-K Item 1.05 each start their own clock on their own trigger. LexLint draws every reporting clock in its corpus on one axis at lexlint.io/clock.

> An incident response that opens by asking which users were affected and where they were is asking a question the record should already have answered, and the hours it takes come out of the same 72.

This is where the two layers meet again. The traceability record is the artefact an incident team reaches for first, and the jurisdiction field is the one that turns it from a description of what happened into a list of who has to be told. A field nobody owns at design time becomes a reconstruction exercise under a statutory deadline.

Scope and limitations

It is not a recommendation that anyone adopt a particular draft or format, and nothing above is required by any single law. The standards efforts named here are not law. What the law requires of a record differs by jurisdiction and is the subject of separate work, published at lexlint.io/analysis/logging. This paper argues only that "which law was in play" has to be answerable from the record, and that answering it needs nothing kept about a person.