ON RETENTION, ERASURE AND THE NUMBER IN THE BOARD PACK

Evidence should be a query, not a project

Every organisation eventually has to prove something about its own past — to a regulator, an insurer, a court, or a board that no longer believes the dashboard. What that costs was decided years earlier by how the platform stores what it saw.

IN BRIEF

Producing evidence is usually a project: someone exports from four systems, reconciles the timestamps, writes a narrative, and hopes nobody asks how a figure was derived. It is expensive because the platforms were built to answer questions in the present. This paper describes what changes when the record of what was observed, judged and done is one object type with time, provenance and retention as first-class properties — and confronts the tension nobody enjoys discussing: the same design that makes evidence retrievable makes personal data retrievable, and an erasure obligation is not satisfied by a filter.

1 · The reconstruction tax

Fourteen months after an incident, a question arrives: what did you know, when did you know it, and what did you do? The team assembles an answer from a closed ticket, a chat thread, an exported log file whose retention window has since rolled, and the memory of an engineer who has left.

The answer produced this way is not necessarily wrong. It is *unverifiable*, which for the purpose at hand is the same thing. Nobody can demonstrate that the narrative matches what the systems recorded, because the narrative was written by a person reading fragments — and the fragments no longer exist in the form they were read.

Three properties of the storage design decide whether this is a week of work or an afternoon: whether the alert kept the events that justified it or a query that must be re-run; whether the actions taken were recorded in the same store as the signals that prompted them; and whether the retention policy was

expressed per record class or per storage tier. The third sounds like an operations detail and is the one that most often destroys the case, because **the events were deleted on a schedule set by cost, not by obligation**.

FIGURE 1 · TWO WAYS TO ANSWER THE SAME QUESTION

RECONSTRUCTED

Four exports · reconciled clocks · a narrative written from fragments · figures nobody can re-derive · six weeks, and the confidence of the person who happened to write it.

RETRIEVED

One store · one clock discipline · the observations, judgements and actions in order · every figure linked to the events that produced it · an afternoon, and the same answer whoever asks.

The difference is not diligence. It is whether the platform treated "what we did about it" as a record of the same kind as "what we saw".

2 · Four properties of a record that will hold up

2.1 · It is stored, not derived

An alert that keeps its evidence says the same thing in two years as it did on the night. An alert that keeps a query says whatever the store returns when the query is re-run — fewer events as retention rolls, different events as sources are reconfigured, and no error to tell you the answer changed. Silent decay is the specific failure, and it is invisible precisely because the interface still renders something.

2.2 · Decisions and actions are events too

A response that ran, an approval that was granted, a legitimacy verdict that was recorded, a rule that was pushed — each is an observation about the estate as much as any process launch. When they share a store and a clock with the telemetry, "what did we do and when" is answerable in the same query as "what happened". When they live in a workflow tool, the two halves must be reconciled by hand, and reconciliation by hand is where dates start disagreeing.

2.3 · Provenance travels with the figure

A number in a report should carry what produced it — the query, the window, the population, the moment it was computed. Without that, a figure is a claim about a computation nobody can repeat, and the first time it is challenged the whole report loses standing. This is also the discipline that catches a stale number before an auditor does: if the provenance is displayed, an old date is visible on the page.

2.4 · Absence is recorded as a fact

"We looked and found nothing" and "we never looked" are different answers, and only one of them is exculpatory. A store that cannot distinguish them cannot be relied on for the most common evidentiary question of all — was this in scope of your monitoring at the time — and the honest platform records coverage and collection gaps as events, so the answer is retrievable rather than argued.

Integrity is a separate claim from retention. Keeping a record is not the same as being able to show it was not altered. Append-only storage, per-record integrity proofs and an independent time source turn "this is our record" into "this is our record and here is why you should believe it has not changed" — a distinction that only matters in the proceedings where everything matters.

3 · The tension: the same design serves erasure and evidence

A platform that retains everything about everyone, indefinitely, so that evidence is always available, is describing a liability rather than a control. Security telemetry is thick with personal data — who signed in from where, which device, which files, at what hour — and the obligations attached to it are not suspended because the purpose is defensive.

The two requirements look opposed and are reconcilable, but only if retention is a property of the record class rather than of the disk it landed on.

Retention is declared per class, with a reason. Authentication events, endpoint telemetry, adjudications and response actions have different lifetimes because they answer different obligations. A single global window is either too short for the audit or too long for the regulator, and usually both.

Erasure is cryptographic where deletion is impractical. Destroying a per-subject key renders that subject's records unreadable across every replica, backup and archive at once — which is the only honest way to satisfy an erasure request against an append-only store, and it is verifiable in a way that a deletion job across seven systems is not.

The erasure itself is an event. What was erased, when, on whose instruction, under which basis — recorded without re-recording the erased content. An erasure with no trace is indistinguishable from data loss, and cannot be evidenced to the person who requested it.

A legal hold suspends the schedule and says so. Holds are exceptions with an owner, a scope and an end date. A hold that quietly becomes permanent is how an organisation ends up retaining, for years, exactly the data it told people it would not keep.

There is a residue this cannot dissolve, and it should be stated rather than engineered around: an incident narrative that names a person's actions remains meaningful evidence and remains personal data. The resolution is procedural, not architectural — minimise what the narrative names, keep the identifying detail in the linked events under their own retention class, and accept that some records are held under an obligation that outlasts a preference. A vendor claiming to have solved this in the product is overselling.

FIGURE 2 · ONE SUBJECT, THREE OBLIGATIONS

`authentication`	`13 months`	supports account-compromise investigation; expires on schedule
`endpoint telemetry`	`90 days`	supports backtesting and hunting; the shortest window that keeps §4 of Paper 04 possible
`actions & approvals`	`7 years`	what the organisation did, and on whose authority — an accountability record
`erasure request`	`key destroyed`	subject's records unreadable everywhere at once; the erasure recorded, the content not

Four lifetimes on one subject, each with a stated reason. A single retention setting cannot express this, which is why estates that have one end up with a policy document that describes something the platform is not doing.

4 · What it changes

The audit stops being a project. Evidence production becomes a query with a date range and a scope, run by the team that already knows the estate, rather than an assembly exercise costing weeks of senior time.

Reporting becomes checkable. When a board figure links to the events behind it, the conversation moves from whether the number is right to what to do about it — and a stale figure is caught by the reader rather than by an auditor.

Privacy commitments become demonstrable. “We keep this for thirteen months and can prove erasure” is a statement the platform substantiates, rather than a sentence in a notice nobody can verify.

Storage cost becomes a governed decision. Per-class retention makes the trade explicit — this window costs this much and buys this capability — instead of a tier setting quietly deciding what the organisation will be able to prove.

5 · Questions worth asking, whoever you are buying from

01 Open an alert from a year ago. Is its evidence stored, or re-queried when I open it?

02 Are response actions and approvals in the same store as the telemetry, on the same clock?

03 Can I set retention per record class with a stated reason, or only per storage tier?

04 How do you satisfy an erasure request against backups and archives — and how is it verified?

05 Does a figure in your reporting carry the query, window and computation time that produced it?

06 Can the platform tell me a source was not being collected during a given window?

07 What integrity proof accompanies a record, and what independent time source does it use?

08 Who inside your organisation can read our events, and is that access itself an event we can see?

6 · Where this sits in the series

The earlier papers argued for enforcement at the edge, alerts that carry their evidence, legitimacy as a distinct judgement, a compilation path from finding to rule, and bounds on what may run unattended. Each generates records — and this paper is about the fact that those records are the organisation's account of itself.

Which reframes the storage decision. It is not an operations concern about volume and cost; it is the decision about what the organisation will be able to prove, to whom, and for how long. That decision is made at design time by people who will not be in the room when it is tested.

> *The retention window is not a storage setting. It is the length of the memory your organisation will have when it is asked to account for itself.*

Helix Counter

ABOUT THIS PAPER

Paper 06 in a series on the architecture of active defense, published by **Helix Counter Research**. Papers 01–05 cover enforcement at the edge, explainable alerts, legitimacy, the compilation path from finding to rule, and bounded autonomy. The series is deliberately vendor-neutral, and the questions in §5 are meant to be asked of us as readily as of anyone else.

Helix Counter holds observations, judgements and actions as one record type with per-class retention, cryptographic erasure, recorded coverage gaps, and figures that carry the query and window that produced them. helixcounter.com