The Signal Ledger — System Design Document

Version: 0.2 (cross-sector expansion) · **Author:** TPOL / Dave Moylan · **Status:** Concept architecture

---

1. Purpose and scope

The Signal Ledger is an institution-agnostic systems design with two coupled objectives. First, to make accountability structural and real-time rather than ceremonial — decisions traceable to their makers, and the people an institution serves able to raise points through a logged, queued, visible channel rather than waiting on a house address session, an annual general meeting, or a complaints process designed to exhaust. Second, to ensure that no genuine concern reported to an institution is ever truly lost: every report is assessed, retained, aggregated, and automatically re-examined when independent signals converge.

Version 0.1 framed the design around government. Version 0.2 generalises it. The failure the Ledger targets is identical wherever an institution holds power over people who must report upward to it: in Parliament, a police force, an NHS trust, a law firm, a bank, a school, and a corporation. The architecture is constant; only the deployment profile changes — who reports, what converges, and which professional body owns the reassessment.

The design does not lower the evidential bar for action in any sector. It changes what happens to information that fails to clear the bar today, so that it remains available to the system tomorrow.

2. Design philosophy: the hindsight problem

The framework rests on one empirical observation. Humans are poor at recognising low-frequency signals in real time but very good at recognising patterns in hindsight — and hindsight arrives after the harm. When a long-running offender is finally caught, the recurring public-inquiry sequence is grimly familiar: investigators are, in the bluntest case discussed in the source conversation, digging up children's bodies from somebody's backyard, and the inquiry asks "were there warning signs earlier?" The answer is usually yes. The harder question — the one this system is engineered to answer — is why those warning signs were not recognised as a pattern at the time.

Inquiry findings across every sector repeatedly identify the same seven barriers: reports

fragmented across agencies; agencies not sharing information; individual reports weak in isolation; no single person holding the complete picture; allegations seeming implausible; limited resources; and assumptions that later proved wrong. None of these is a failure of information. All of them are failures of connecting, retaining, or recognising it. The Signal Ledger therefore targets the connective tissue, not the collection.

The economic argument is an asymmetry: retention is cheap; hindsight is catastrophic. The cost of preserving a genuine weak signal for future analysis is trivially small beside the cost of ignoring one.

3. The universal pattern

Every large institution, regardless of sector, processes concerns through the same legacy lifecycle:

```
Report → Assessment → Closure
```

If the report cannot be substantiated alone, the file closes and the information functionally dies. Police complaints, patient concerns, client grievances, mis-selling complaints, safeguarding worries, HR cases — the form differs, the lifecycle does not. The Signal Ledger replaces it everywhere with:

```
Report → Assessment → Retention → Aggregation → Pattern Detection → Reassessment
```

Closure of a case is decoupled from disposal of a signal. This single decoupling is the whole framework; everything else in this document is the machinery required to make it safe, lawful, and operable at national scale.

4. System A — The Accountability Layer

4.1 Decision Ledger

An append-only record of significant decisions capturing: the decision, the decision-maker(s), the timestamp, the evidence considered, the alternatives rejected, and the stated rationale. In government this covers ministerial and departmental decisions; in a hospital trust, board and clinical-governance decisions; in a corporation, executive decisions above a defined materiality threshold. The design goal is constant: “who is responsible?” becomes a lookup, not an investigation. Records are immutable once written; corrections supersede, never delete, preserving the full history of what was claimed and when.

4.2 Stakeholder Channel

The constituent-to-PM channel of v0.1, generalised. Every institution defines its stakeholders — constituents, patients, clients, customers, employees, parents — and provides a real-time submission route that is logged with a unique reference, queued, visible (with submitter identity protected by default), and answered on the record. The channel makes two promises only: every point enters the ledger, and the aggregate pattern of what stakeholders are raising becomes visible, queryable data. Ten thousand independent stakeholders raising the same issue is itself a convergence event, routed to the same pattern engine as System B.

4.3 Reciprocity Audit

The institution applies to itself the standards it applies to those it serves. Response deadlines are tracked in the Decision Ledger; missed deadlines are logged events. Access to any ledger is logged in that ledger. The perceived gap between stated principles and actual practice — the root of most public distrust — is closed architecturally rather than rhetorically.

5. System B — The Signal Layer

5.1 Components

Intake. Low-friction reporting with structured capture of entities (persons, locations, organisations), described method or behaviour, timeframe, and the reporter's relationship to events. The division of labour is firm in every sector: citizens, patients, clients, and employees report concerns; professionals investigate them. Intake never asks, encourages, or rewards amateur investigation.

Triage and assessment. Standard professional assessment, unchanged from current sector practice. The outcome (actioned, referred, unactionable, implausible-as-stated) is recorded as metadata — including the assumptions on which the assessment rested, because those assumptions are precisely what later convergence may justify revisiting.

Retention Ledger. Append-only, access-logged storage of every report and its assessment, with statutory retention horizons set per category — longest where the hindsight cost is highest (safeguarding, patient safety, abuse of power).

Pattern Engine. Continuous entity resolution and clustering across the retained corpus, seeking convergence along four dimensions: same individuals, same locations, same methods, same or overlapping timeframes. The core weighting principle is **independence**: reports are scored by provenance separation — different reporters, no shared social or informational pathway, different intake routes, different times. Five unconnected reports describing the same behaviour outweigh fifty copies of a single rumour. Raw volume without

independence scores low; modest volume with high independence scores high.

Convergence triggers. Threshold rules that promote a cluster from background to review — illustratively, a new report matching two or more prior independent reports on at least two dimensions generates a reassessment task. Thresholds are tunable per risk category: safeguarding and patient-safety clusters trigger at lower convergence than billing complaints.

Reassessment workflow. A triggered cluster routes to a human professional with the full cluster in view — the complete picture that, in the legacy model, no one person ever held. The reviewer's question is not "is each report proven?" but "do the assumptions made when each report was assessed in isolation survive contact with the cluster?"

5.2 Worked example: the four reports

A report arrives. Viewed alone it is unusual, unlikely, poorly evidenced, and difficult to verify. Legacy outcome: assessed and closed. Ledger outcome: assessed, marked unactionable, retained. Years later — or the same week — the pattern engine finds that multiple people have reported similar behaviour; the reports involve the same individuals; the same locations; similar methods; and emerged independently of one another. A convergence trigger fires.

The load-bearing distinction: four independent reports do not prove the allegation. They justify revisiting the assumptions made when each report was viewed in isolation. The system changes when professionals look, not what counts as proof. Due process, evidential standards, and investigative rigour remain fully intact downstream.

5.3 Worked example: how investigations break open

Many major historical investigations follow one sequence: a new report arrives; analysts review historical records; similar past reports are discovered; connections become visible; what looked isolated begins to look systemic. The Signal Ledger does not invent this process — it automates the archive review, so the connection-finding step happens by design rather than depending on one diligent analyst thinking to look back.

5.4 Worked example: the escalation ladder

Information gains meaning through context. One report is a report. Five similar, independent reports may be a pattern. Fifty similar reports may indicate a systemic issue. The ladder is a heuristic, not an algorithm — independence weighting operationalises it. The governing maxim of the whole layer: **today's dismissed anomaly may be tomorrow's confirmed pattern.**

6. Sector deployment profiles

The architecture is constant. Each profile specifies four variables: signal sources, primary convergence dimensions, reassessment owner, and regulator interface.

6.1 Government and politics

Signal sources: the Stakeholder Channel (constituent submissions), departmental complaints, whistleblowing routes. Convergence: policy area, department, named decision-maker, locality. Reassessment owner: parliamentary committees and the Ledger Authority. The Decision Ledger carries the heaviest load here: ministerial decisions, evidence considered, alternatives rejected — published, immutable, queryable. Constituents raise points in real time; the queue is public; the pattern of what a nation is asking its government becomes data the government cannot un-see.

6.2 Policing

Signal sources: public complaints, internal professional-standards reports, vetting data, court outcomes. Convergence: officer identity, behaviour type, force, location. The sector's defining failure mode is **fragmentation by transfer** — an officer moves forces and the record loses continuity. Deployment requires a national retention ledger spanning all forces, so that seven weak complaints across three forces over a decade surface as one convergence event with the complete picture attached. Reassessment owner: professional standards departments with independent oversight (in the UK, the IOPC interface). Vetting becomes a continuous query against the full national corpus rather than a point-in-time snapshot.

6.3 Healthcare and the NHS

Signal sources: patient complaints, family concerns, staff whistleblowing (Freedom to Speak Up routes), clinical incident reports, mortality and outcome statistics. Convergence: clinician identity, ward/unit, clinical method, temporal clustering of incidents. The sector's defining failure mode is **silo separation** — the complaints office, the incident system, the HR file, and the mortality data never meet, so a clinician whose incidents cluster across wards and trusts stays invisible. Deployment federates these silos into one pattern engine per trust, federated nationally. Reassessment owner: clinical governance, with CQC and professional-regulator interfaces (GMC, NMC). The design is reciprocal: a patient maintaining a structured master timeline of their own treatment, correspondence, and unanswered questions is operating a personal signal ledger against the institution, and the institution's ledger should be required to answer to it.

6.4 The legal sector

Signal sources: client complaints to firms, regulator complaints (SRA), negligence claims, costs disputes. Convergence: practitioner identity, conduct type, claimant independence. The sector's defining failure mode is **contractual forgetting** — the confidential settlement. A practitioner can leave a trail of quietly settled claims, each sealed by an NDA, each invisible to the next client and frequently to the regulator. The Ledger retains the fact, category, and shape of each settled signal even where content is sealed; confidentiality binds the parties, not the regulator's pattern engine. Five independent claimants alleging the same conduct type against the same practitioner is a cluster the SRA sees regardless of what each settlement says. Settlement ends a dispute; it no longer erases a signal.

6.5 Financial services

Signal sources: customer complaints, ombudsman referrals, internal whistleblowing, conduct-risk reporting. Convergence: product, firm, sales method, branch/desk. The sector has already demonstrated the failure mode at national scale: mis-selling scandals are, structurally, millions of weak signals assessed in isolation for years — each complaint individually deniable, the systemic practice invisible until the aggregate forced itself into view through ombudsman backlogs and litigation. A Ledger clustering complaints by product and method surfaces the convergence in months, not decades. Reassessment owner: conduct-risk functions with FCA interface. Whistleblowing gains the same benefit: the lone insider report acquires its true weight beside the retained reports of insiders who came before.

6.6 Education and safeguarding

Signal sources: pupil and parent concerns, staff reports, low-level concerns logs, referral records. Convergence: staff identity, behaviour type, institution sequence. The sector's defining failure mode has a name — **passing the trash**: a staff member departs under unproven concerns, the reference stays neutral, and the history launders itself with each move. The Ledger retains low-level concerns as signals (explicitly not accusations) visible to designated safeguarding professionals across institutional boundaries, so three quiet departures from three schools surface as one pattern. Reassessment owner: designated safeguarding leads and the LADO function, with DBS and Ofsted interfaces. This profile carries the strictest access controls and the longest retention horizons in the entire framework, because the hindsight cost is measured in children.

6.7 Employers and large organisations

Signal sources: HR complaints, grievance and disciplinary records, exit interviews, anonymous ethics lines. Convergence: subject identity, conduct type, team/manager aggregation. Defining failure modes: **investigator amnesia** (each new complaint investigated from zero because the prior unsubstantiated file is closed and forgotten) and

discarded exit data (the richest signal source most organisations possess, collected and binned). The corporate Ledger retains both: complaints clustered by subject and conduct type across years and restructures; exit-interview themes aggregated by team. The serial problem protected by settlements and staff churn becomes a query result instead of an open secret. Reassessment owner: independent ethics or audit functions reporting to the board, not to the management chain being measured.

6.8 Safety-critical industries: the proof the model works

Aviation has operated the Signal Ledger's core logic for decades. Confidential incident and near-miss reporting systems retain every report — trivial, unproven, embarrassing — aggregate them across operators and borders, and treat convergence as the trigger for action. No single near-miss proves a design flaw; thirty independent ones ground a fleet. Aviation became the safest complex industry on earth not by collecting more information than healthcare or policing, but by refusing to forget any of it, and by separating safety reporting from blame so that signals flow freely. Two transferable lessons: retention plus aggregation works at global scale, and a just-culture intake (report without fear) multiplies signal volume precisely where it matters. Rail, nuclear, and pharmaceutical pharmacovigilance run variants of the same machinery. The framework's challenge to every other sector: why are your weak signals treated with less respect than an airline treats a bird strike?

7. The adversary: mechanisms of institutional forgetting

A cross-sector design must name what it is up against. The Ledger's true adversary is not malice but the ordinary machinery by which institutions forget:

confidential settlements and NDAs (legal, corporate); fragmentation by transfer (policing) and by silo (NHS); reference laundering (education); investigator turnover and case-by-case amnesia (HR); records-retention minimisation policies that delete precisely the unproven material that pattern detection needs; and organisational restructures that orphan historical data. Each sector profile in §6 exists to counter that sector's dominant forgetting mechanism. A deployment that does not disable the local forgetting mechanism is decoration.

8. Scale architecture

National, multi-sector deployment implies volumes in the tens of millions of signals. The load-bearing technical problems are:

Entity resolution at scale. The same individual appearing across forces, trusts, schools, and employers must be resolvable without creating a universal surveillance identifier. The

design uses sector-scoped pseudonymous identifiers with a cryptographic linking function exercisable only by the Ledger Authority under logged, purpose-coded authority — linkage is an audited event, not an ambient capability.

Independence scoring. Provenance-separation metrics (reporter distinctness, route distinctness, temporal spread, absence of shared informational pathway) computed at ingest and recomputed as clusters grow. Calibrated against historical inquiry datasets where ground truth is known.

Federation before unification. Each institution runs its own ledger; cross-institution pattern detection operates over federated queries with sector-appropriate access law. No central pool of raw narratives — the pattern engine moves to the data, not the data to a central honeypot.

Trigger economics. Reassessment capacity is finite. Trigger thresholds are tuned so that the professional review queue is fed at a rate the sector can staff, with risk-weighted prioritisation — and the firing rate, queue depth, and outcome statistics are published, so under-resourcing the review function is itself visible in the Decision Ledger.

9. Data model sketch

```
SIGNAL        { id, sector, institution_id, received_at, intake_route,
                reporter_ref (protected), entities[], locations[],
                methods[], timeframe, narrative, category }
ASSESSMENT    { signal_id, assessor_ref, outcome, assumptions[],
                assessed_at }
SETTLEMENT    { signal_id, exists: true, category, date,
                content: sealed }        // legal-sector forgetting counter
CLUSTER       { id, signal_ids[], dimensions_matched[],
                independence_score, sector_scope, status }
TRIGGER_EVENT { cluster_id, rule_fired, created_task_ref, fired_at }
DECISION      { id, institution_id, decision, makers[], evidence_refs[],
                alternatives[], rationale, timestamp }
ACCESS_LOG    { actor_ref, record_ref, purpose_code, at }   // recursive
LINKAGE_EVENT { pseudonym_a, pseudonym_b, authority_ref,
                purpose_code, at }       // cross-sector entity resolution
```

All stores are append-only; supersession, never deletion, with lawful-erasure handled by cryptographic redaction that preserves the existence and shape of the record.

10. Threat model and safeguards

Retention is not accusation. A stored signal carries no presumption of guilt against any

named person, in any sector. Cluster status is an internal review state, never a public designation. This is doubly critical in the education and policing profiles, where a retained-signals regime touching individual reputations must be matched by a statutory right of reference, review, and challenge for named subjects.

Weaponised reporting. Coordinated false reporting is the principal attack in every sector — disgruntled litigants, vexatious complainants, organised campaigns. Defences: independence weighting (coordinated reports share provenance pathways and score low); orchestration detection (the convergence machinery that finds genuine patterns equally finds campaigns — identical phrasing, temporal bunching, shared origin); and consequences routed through existing false-reporting law. The archive cuts both ways by design.

Privacy and proportionality. Reporter identity protected by default; subject data access-controlled by sector, category, and purpose; every access logged recursively, with the watcher's query itself subject to pattern detection. Retention horizons and access purposes set in statute per sector, audited by an independent body publishing into the Decision Ledger. Cross-sector linkage is the highest-risk operation in the system and is correspondingly the most tightly gated (§8).

False-pattern risk. Convergence triggers create review tasks, not actions. Humans own every consequential judgement. Trigger thresholds, firing rates, and reassessment outcomes are tracked and published, so the system's own error rate is measurable — reciprocity applied to the machinery.

Surveillance creep. Scope is bounded to reports voluntarily made and decisions formally taken. The system ingests what stakeholders and institutions put into it; it does not collect ambient data. Any scope expansion requires a Decision Ledger entry naming who expanded it and why — expansion by stealth is structurally impossible without leaving its own audit trail.

Chilling effects. A just-culture intake, proven in aviation, is mandatory in every profile: reporting in good faith carries protection, and safety/conduct signals are firewalled from routine performance management so that the ledger increases reporting rather than suppressing it.

11. Governance

An independent Ledger Authority — separate from every institution whose performance the ledgers expose — owns the infrastructure, the trigger rulebook, the linkage function, and the audit programme. Sector regulators (illustratively, in the UK: IOPC for policing; CQC, GMC and NMC for health; SRA for legal; FCA for finance; Ofsted, DBS and LADO functions for education; HSE for workplace safety) own reassessment standards within their domains

and receive cluster escalations under statutory gateways. The Authority's own decisions live in the Decision Ledger; its access lives in the Access Log; its performance is published. Statute defines retention horizons, access purposes, subjects' rights of reference and challenge, reporters' rights to track their own submission, and the reciprocity obligations on every participating institution.

12. Phased implementation

Phase 1 — Decision Ledger pilots. One government department, one NHS trust, one corporation publish decisions in ledger format. Lowest risk, immediate transparency value, builds the append-only infrastructure and the audit culture.

Phase 2 — Stakeholder Channels. Real-time submission, public queues, topic clustering — government constituent channel first, then patient and customer channels. Proves intake and aggregation on low-sensitivity data.

Phase 3 — Signal Layer in the highest-hindsight-cost domains. Education safeguarding and one NHS patient-safety domain, where legacy failure modes are best documented by existing inquiries and the moral case is strongest. Strictest access controls debugged here first.

Phase 4 — Policing national retention ledger. The fragmentation-by-transfer problem solved with mature components and a tested rights-of-challenge regime.

Phase 5 — Regulated professions and finance. Legal-sector settlement registration and financial-conduct clustering, riding on regulator gateways established in earlier phases.

Phase 6 — Cross-sector federation. Deliberately last. Federation of mature ledgers is tractable; federation of immature ones reproduces the fragmentation it was meant to fix — and the linkage function, the system's most powerful and most dangerous capability, is switched on only when every safeguard around it has years of audit history.

13. The seven principles

1. Accountability should be structural, not dependent on trust.
2. Archive rather than forget.
3. Aggregate rather than isolate.
4. Look for convergence rather than certainty.
5. Citizens report; professionals investigate.
6. Institutions answer to the standards they impose — reciprocal accountability.

7. Build feedback loops so that hindsight is never the only teacher.

14. Open questions

How are independence scores calibrated against real inquiry datasets, sector by sector? What does the statutory right of reference and challenge look like for a person named in retained signals they have never been told exist — and when does notification itself become the harm? Where exactly do retention horizons sit per category and sector? How does settlement registration survive contact with legal-sector lobbying? How is the just-culture firewall enforced when the same institution runs both the ledger intake and performance management? How does the Stakeholder Channel resist becoming a popularity engine rather than a signal source? These are flagged, not solved — a proper design names its unknowns.

A society learns more effectively when information accumulates, patterns are recognised, and accountability remains traceable across time — in every institution, not just the ones we have already caught failing.