

BOUNDARY BLINDNESS IN INSURANCE

*Early Findings and a
Call for Industry Validation*

Alexander D. Barrett

SENIOR FELLOW

VERIDA CHARTER FOUNDATION

A NONPROFIT RESEARCH FOUNDATION ON
CROSS-FIRM COORDINATION IN REGULATED
MULTI-PARTY MARKETS

JUNE 2026



Boundary Blindness in Insurance

*Early Findings from Eight Months of Senior-Practitioner Research,
and an Invitation to Phase Four*

Verida Charter Foundation

*A nonprofit research foundation on cross-firm coordination
in regulated multi-party markets, 501(c)(6)*

Alexander D. Barrett, Senior Fellow

June 1, 2026

Estimated reading time: 55-75 minutes for the full paper; 10 minutes for the executive summary.

This paper presents early findings from the first eight months of a twelve-month research programme running from October 2025 to October 2026. Phases one through three (practitioner interviews, convergence-cluster analysis, and triangulation against published literature) are now complete; Phase Four — set out in Section 9 — occupies the remaining four months and is structured to test the findings empirically against operating environments.

The full aetiology — methodology, regulatory analysis, evidence base, and symptom-to-cause mapping — is set out in an underlying thesis research paper, *The Aetiology of Boundary Blindness*, a working thesis available on request. The findings here are converging, not concluded; Phase Four will test them against three platform operating environments under a four-week sandbox protocol with no commercial commitment.

Practitioner quotations are anonymised at role and tenure level. Public authors of published artefacts are named directly. The Verida Charter Foundation is a nonprofit research foundation focused on cross-firm coordination infrastructure in regulated multi-party markets.

The Case in Brief

Insurance does not merely have a technology problem, a claims problem, an AI governance problem, or a regulatory-evidence problem. Those are symptoms of a deeper structural failure: decision evidence does not travel reliably across firm, platform, and jurisdictional boundaries. The paper calls this *boundary blindness*. The market has normalised workarounds so thoroughly that it no longer sees the common cause.

This is not a people problem. It is an architecture problem. Insurance has built *rails* for policy data, claims data, bordereaux, documents, messages, and settlement records. But it has not built rails for the decision context. The actual decision trail — rationale, authority, evidence, judgment, and acceptance — remains trapped in emails, local systems, fragmented platform logs, and practitioner memory. That is why followers cannot always see what the lead reviewed, audits take weeks, and old decisions become hard to defend.

The paper identifies nine symptoms of the same missing layer: followers cannot see what leads reviewed; leaders preserve narrative control; selective disclosure is required; decision-state must hold at execution; decisions are hard to reconstruct years later; underwriting, renewal, and claims teams lack structured handover; institutional memory is leaving the workforce; AI controls sit in policy rather than execution; and AI decision provenance is missing. These are not separate problems. They are one structural gap appearing in different settings.

The proposed answer is Portable Trust Infrastructure, or PTI. PTI is not a new claims platform, a blockchain pitch, a centralised data pool, or a workflow system. It is a coordination layer that sits beside existing systems and creates a signed, portable decision record: an acceptance signal. That signal records that a person or system reviewed specific evidence under defined authority, in accordance with relevant rules, and reached a conclusion at a particular time. The underlying data does not need to move. The proof of the decision can.

That matters because insurance depends on controlled asymmetry. Leads, followers, brokers, reinsurers, MGAs, TPAs, and auditors do not all need to see everything. Often, they should not. But they need sufficient verifiable evidence to safely rely on their decisions. PTI preserves selective disclosure: followers can verify that decision work was done without reading private notes; auditors can see what they are entitled to see; counterparties do not receive blanket access.

The urgency comes from three forces converging at once: institutional memory is disappearing, regulators are demanding evidence of intent and oversight, and AI is accelerating decisions across boundaries already under strain.

The paper is honest about what remains unproven. The diagnosis is converging, not final. It is anchored mainly in UK and US insurance, especially Lloyd's and specialty markets. It does not yet provide a complete protocol, security model, or tested interoperability layer. The call to action is Phase Four: test the acceptance signal in real platform environments and move from persuasive diagnosis to empirical proof.

The core message is simple: insurance has spent thirty years digitising documents and transactions, but not decisions. That was tolerable while trust, memory, and relationships absorbed the gap. Under AI, regulation, workforce transition, and complex multi-party claims, it is no longer tolerable. The next layer the industry needs is portable, verifiable decision evidence.

Executive Summary

Insurance has a coordination problem hidden in plain sight. Across the global property and casualty market, between 2% and 5% of claims by volume drive about 85% of indemnity exposure:¹ the multi-policy, multi-party, multi-jurisdictional losses that most depend on cross-firm coordination, and where the absence of portable '*decision evidence*' is most consequential. At the volume end of the market, the same gap surfaces through AI-mediated decisions at scale, outsourced TPA arrangements under regulatory scrutiny, and Consumer Duty's per-decision evidence requirements across personal lines and small commercial. This paper argues these exposures are not separate problems.

This research argues it is the surface of a single structural condition: the inability of *decision evidence* to travel reliably across firm, platform, and jurisdictional boundaries. The paper names the industry's inability to see that the nine symptoms below all appear to trace to this one underlying condition, and to act on it: **boundary blindness**.

The condition has been operationally tolerable for three decades because human institutional memory absorbed the cost. Three concurrent forces are now removing that absorber simultaneously. The window for architectural intervention is approximately 18 to 24 months.

A Group Executive Chairman of a national loss-adjusting firm, asked whether his firm's reports were tracked after they leave the adjuster's hands, was direct:

*There is no control of who it goes to at all.*²

Over eight months from October 2025 to mid-May 2026, this research conducted 61 semi-structured interviews with senior practitioners (47 within the United Kingdom and United States insurance markets³, and 14 across adjacent regulated multi-party domains). It analysed 21 independently published sources from regulators, reinsurers, trade bodies, and adjacent sectors, and triangulated findings against contemporary trade-body publications and institutional research.⁴

The research is anchored in UK London Market practice, the cohort with the deepest representation, but the diagnosis is structural. It applies wherever cross-firm coordination between regulated

¹ The specific 2–5% / 85% concentration figure is attributed to Adam Humphrey, MD Complex Claims International and Immediate Past President of the Chartered Institute of Loss Adjusters (CILA), in primary-research interview for this paper, May 2026. Humphrey's twenty-five-year career adjusting major complex claims across the international placement chain anchors the practitioner synthesis. The general phenomenon — Pareto-distributed severity in which a small share of large losses drives a disproportionate share of total indemnity exposure — is well-established in the actuarial literature (Bahnmann, Distributions for Actuaries, CAS Monograph Series No. 2). The most recent industry illustration is Swiss Re Institute, sigma 1/2025, finding that secondary perils accounted for 92% of the USD 107 billion in global insured catastrophe losses in 2025; the same heavy-tail shape at the per-claim level produces Humphrey's figure.

² Primary-research interview with the Group Executive Chairman of a United Kingdom national loss-adjusting firm carrying over four decades of executive experience in the loss-adjusting sector. One of 61 senior-practitioner interviews conducted between September 2025 and May 2026 for this research. Notes on file with the author. Interviewee anonymised at source; identifying details (firm name, exact role) omitted to preserve confidentiality.

³ This paper synthesises the empirical and analytical findings of phases one through three of an ongoing research programme. The methodology section sets out the interview design, anonymisation discipline, and triangulation framework. Full methodological detail is in the underlying research paper, The Aetiology of Boundary Blindness, working thesis, Dec. 2026.

⁴ The twenty-one sources are catalogued in Section 5. Full bibliographic details are in the References list. The convergence-cluster identification in this paper is the first known compilation of these sources as a single analytical pattern.

insurance participants occurs: US P&C subscription business and reinsurance chains, MGA delegations, TPA arrangements, personal-lines volume operations, and cross-jurisdictional placement. The architectural answer is correspondingly market-agnostic.

The paper proposes a four-part diagnostic structure:

1. **Nine symptoms** are present in the operational reality of insurance every day. They surface as routine overhead, such as rework, escalation volume, audit preparation time, repeated meetings, email threads, and callbacks.
2. **One condition** is consistent with all nine: a structurally missing decision-evidence layer. The causal claim awaits Phase Four; the convergence is documented now.
3. **One architectural intervention** would address the condition: Portable Trust Infrastructure (PTI), centred on the primitive that the paper calls the *acceptance signal*. Both are specified in Section 6.
4. **Boundary blindness** names the condition this research argues has prevented the industry from recognising the nine symptoms as manifestations of a common underlying pattern. Practitioners normalise the workarounds; the workarounds become invisible. The work continues; the cost is absorbed elsewhere.

The condition appears at three boundary types that compound: cross-firm (between any two regulated insurance participants such as insurers, reinsurers, brokers, MGAs, TPAs, and loss adjusters, in subscription-market lead-follow, broker placement chains, MGA delegations, TPA arrangements, and reinsurance flows), cross-platform (between distinct software systems, often within a single firm), and cross-jurisdictional (between distinct regulatory regimes). Each is unpacked in Section 1; each appears in the nine symptoms catalogued in Sections 2 to 4.

Why now: three concurrent forces are removing the absorber, the human institutional memory that has compensated for the missing decision-evidence layer for three decades. A claims-platforms consultant with three decades of London Market experience named the regulatory pivot:

*You're going to have to prove not only you did it, but that you did it with intent.*⁵

The shift from action-evidence to intent-evidence is the load-bearing change. Action-evidence is what the current infrastructure can already produce; intent-evidence requires a decision context that does not survive transit across firm boundaries. The Fifth Hurdle, FCA delegated authority expansion, Solvency II exit plans, DORA, the EU AI Act, APRA on AI, and the NAIC's AI Systems Evaluation Tool pilot each demand the latter; none allows the substitution of the former.

The other two forces compound it. The workforce demographic transition is removing the human memory that absorbed boundary blindness for three decades: the under-thirty London Market cohort is falling from 24% to a projected 7%. Artificial intelligence is multiplying the load that absorbers can no longer carry, scaling into claims and underwriting workflows at machine speed across boundaries

⁵ Primary-research interview with a claims-platforms consultant carrying three decades of London Market experience across multiple major platform vendors. One of 61 senior-practitioner interviews conducted between September 2025 and May 2026 for this research. Notes on file with the author. Interviewee anonymised at source; identifying details (firm, exact role) omitted to preserve confidentiality.

that human-paced workarounds were not designed to cross. The convergence-cluster sources distil it cleanly:

*"The signal is not the decision. The decision is only real when it holds under consequence."*⁶

The convergence cluster. Twenty-one independently published sources across six jurisdictions and seven sectors have identified variants of the same architectural failure, twelve of them within the past eight months. Anchor voices include APRA (Australia), Munich Re (Germany), the Geneva Association (Switzerland), the NAIC (United States), and a London Market four-corner alignment across LIIBA, LMA, Lloyd's, and the IUA (United Kingdom), with cross-vertical reinforcement from manufacturing and industrial AI. The full set is catalogued on chart p. 21. None of the twenty-one references another's work; the root cause is invisible from any single vantage and appears only when the sources are read as a single pattern. The boundary blindness this paper diagnoses applies to the cluster authors as much as to the practitioners they study.

Anticipating the misreads. The architecture proposed here produces audit evidence, not decision surveillance. Anchored in the Insurance Act framework, which already requires insurers to demonstrate how and why decisions were reached, the structured, signed record serves as a defensibility tool rather than an oversight imposition. Controlled asymmetry is preserved by design: selective disclosure allows followers to verify that decision work was performed to a documented standard without reading the lead's working notes, and the same logic operates at the broker layer. Lloyd's Delegated Data Manager (DDM) has already operationalised this primitive within a single market; Portable Trust Infrastructure (PTI, proposed later in Section 6) extends it across the cross-firm, cross-platform, and cross-jurisdictional surface.

The research has reached the point where empirical validation against operating environments is the appropriate next step. Phase Four tests the architectural intervention with three platform participants over twelve weeks each, with no commercial commitment.

Phase Four — Empirical Validation

Three platform participants across the UK, Europe, and, where practical, the US and Australian insurance domains. Each runs for one calendar quarter (about 12 weeks of structured engagement), with three-week staggered starts. About 25–30 hours per platform across the quarter; no purchase or procurement commitment. Engagements complete between October and December 2026; findings published Q1 2027. The research benefits from any outcome, including a negative one. Indication of interest by 26 August 2026. Full structure in Section 9.

Phase Four is the empirical step that distinguishes a working diagnosis from a tested one. It is designed to validate whether the acceptance signal can be produced and consumed across firm and platform boundaries in real environments, and to surface failure modes. The convergence-cluster, demographic, and regulatory analyses rest on existing evidence and are not the subject of Phase Four testing.

⁶ Perlstein, Brian, and Michael Carroll (2026), *The Signal Is Not the Decision*. Long-form essay, 5 May 2026. The phrase is also the essay's title.

Author disclosure

This research was self-funded under the Verida Charter Foundation (a 501(c)(6) US nonprofit); no party funded or directed it. To enable Phase Four validation, the author built prototype implementations of the concepts described here; these prototypes are the author's personal intellectual property. The VCF Protocol remains open under Apache 2.0, and any party may build a compatible implementation independently. The prototype backend may later be operated by a successor commercial entity — this affects neither the Protocol's openness nor any platform's ownership of its connector kit. Findings will be published in full, including any unfavourable to the prototype work. Should Phase Four engage sponsors, the foundation's editorial-independence policy ensures funder relationships do not influence findings; full terms provided on engagement.

What this paper is — and what it is not

The paper presents a working diagnosis, a proposed architectural intervention, and an empirical test (Phase Four). It is not a critique of existing platforms; not a proposal for centralised data aggregation; not a blockchain proposition; not a new claims platform competing with existing ones; not a workflow optimisation programme. The research is agnostic about who should build the missing layer. What it proposes is that the architectural diagnosis is increasingly supported by both practitioner evidence and the published literature, and that an open specification, validated by the platforms whose customers experience boundary blindness most acutely, is the appropriate next step.

Terms used in this paper

Several terms recur throughout the paper:

Boundary blindness — the industry's structural inability to see that its many operational symptoms appear to share a single underlying condition: the missing decision-evidence (context) layer.

Decision context/evidence — who reviewed what, under what authority, applying what evidence, reaching what conclusion, and the age of the decision.

Portable Trust Infrastructure (PTI) — the architectural category proposed by this paper: event-driven infrastructure that sits alongside existing software platforms and produces signed, portable decision records that travel across firm, platform, and jurisdictional boundaries without touching the underlying data. The acceptance signal is the primitive PTI that the sigillum seals. PTI is a coordination layer, not a new platform or centralised data aggregation.

The acceptance signal — a structured, signed, portable decision record that travels independently of the data it attests to. Carries: who reviewed what, under what authority, applying what evidence, applying which rules or guidelines, reaching what conclusion, and at what point in time. Verifiable across firm, platform, and jurisdictional boundaries without requiring access to the underlying source material.

The Sigillum — historically, the wax seal that authenticated and protected a medieval document. In the architecture proposed in this paper, the sigillum is represented by the cryptographic layer that proves an acceptance signal has not been altered, governs who may read or append to it, and preserves the integrity of the decision across the boundaries it must cross.

Controlled asymmetry — the deliberate information asymmetry between lead and follower (or broker and reinsurer) that the market has evolved for legitimate operational reasons. Any architectural answer must respect it, not override it.

Selective disclosure — the architectural property that allows different recipients to access different subsets of a single record, at the document, field, or region level.

Data rails — the rules, schemas, and protocols that set the standards for how data is structured and exchanged between firms in the subscription market. In insurance: ACORD message standards, the specifications behind ECF and PPL, and emerging standards for AI-extracted data. The rails define the format; they do not carry the decision context attached to the data.

How to read this paper

The paper is structured to serve four reader audiences in parallel:

1. **Claims and operational leaders** may focus on Sections 3–4, which catalogue the nine operational symptoms documented in practitioner interviews and the regulatory pressures that have elevated them.
2. **Platform and architecture teams** may focus on Sections 5–6, which describe the convergence cluster and the architectural implications — including the form the acceptance signal must take to cross firm, platform, and jurisdictional boundaries simultaneously.
3. **Regulators, supervisors, and governance stakeholders** may focus on Sections 3, 5, and 7, which describe the symptoms, the institutional convergence of recognition, and the forcing functions that distinguish the present from the past three decades.
4. **Senior executives and board-level readers** may focus on Sections 1, 7, and 9, which present the diagnostic framework, the why-now case for architectural intervention, and the Phase Four invitation for empirical engagement.

1. The Diagnostic: Boundary Blindness

The opening question of this research, in October 2025, was whether the cross-firm coordination problems described informally in the London Market for decades had a single shared structural cause. The answer, accumulated across sixty-one interviews and the published literature, increasingly points to one: the condition is not new; the cost of leaving it structurally unaddressed compounds under regulatory pressure faster than the institutional response; and artificial intelligence, scaling into operational use through 2026 and beyond, will compound the same cause at machine speed.

The three-decade frame this paper uses is anchored to the post-1996 Lloyd's market that emerged after Reconstruction and Renewal, when the subscription market scaled beyond its collocated, handshake-based origins.⁷ The data rails — the standards, schemas, and protocols defining how data

⁷ Lloyd's Reconstruction and Renewal completed on 4 September 1996, when the Equitas reinsurance vehicle assumed all pre-1993 liabilities and the modern institutional Lloyd's market began operating in its current form. The frame is internally consistent across the paper: the senior practitioners with "three decades" of market experience referenced in the interview cohort are the

moves between firms — were built over the intervening period: ACORD message standards, ECF for claims, broker placement platforms, and, more recently, AI-enabled extraction. The *decision-evidence* layer was not. Institutional memory has absorbed the gap for a generation, and the generation now leaving the workforce carries the memory with them.

Practitioners do not describe their working environments as breaking. They describe them as normal. Several of the workarounds (manual re-keying between platforms, the demographic transition in long-tenured staff, duplicated verification across firms) are discussed constantly in the trade press and named in industry research. What remains invisible is their shared structural cause: the workarounds have accumulated over the past three decades as the conditions of the work have been treated as separate operational realities rather than as manifestations of a single underlying coordination deficit.

The Group Executive Chairman of a national loss-adjusting firm, with over four decades of executive experience in the United Kingdom, described what happens to his firm's reports after they leave the adjuster's hands. The report goes to the instructing party (a Lloyd's syndicate, a managing general agent, or a company-market insurer), who may then pass it on to lawyers, brokers, assessors, or sometimes the insured. The adjuster has no visibility of where the report goes after dispatch or who acts upon it. Asked directly whether this reuse was seen and tracked, he confirmed:

There is no control of who it goes to at all.

The Chairman's point is operational, not legal. Data protection law (the UK GDPR and the Data Protection Act 2018) restricts who the report can lawfully be shared with; carrier claims management systems and the Electronic Claims Files capture some downstream activity within the formal claims workflow. What is absent is operational visibility: onward distribution to lawyers, retro-reinsurers, brokers' follow markets, and other recipients typically falls outside that tracking, and the originator receives no signal of which legally permitted recipients the report actually reaches or how it is being reused downstream.

Asked whether the gap worsens when claims cross platform boundaries, the same Chairman confirmed:

There is a problem across platforms... we're not integrated... there are links between systems and different, but generally have that breakdown. And sometimes you will find where an auditor will look at their file without looking at how far, and we will kind of go, no, hang on, you haven't got all the information.

Claim references, document classifications, line-of-business identifiers, and jurisdiction codes move with increasing efficiency via ACORD-formatted messages, ECF, broker workflows, and platform integrations. The *decision context* (who reviewed what, under what authority, with what conclusion, applying which rules, and whether the decision remains current) does not. The *acceptance signal* this paper proposes attaches to that decision context, not to the data; the data continues to move as it

cohort whose careers span this same period, and who have carried the institutional memory the paper identifies as the absorbing mechanism. Practitioners with "four decades" of experience referenced in the paper are those whose careers also include the pre-1996, pre-R&R market.

does today, and a signed, portable record of the context above accompanies each material decision made on its basis.

The cost of the structural deficit is not evenly distributed across the claims population. Adam Humphrey, MD Complex Claims International and Immediate Past President of the Chartered Institute of Loss Adjusters (CILA), with twenty-five years of adjusting major complex claims, named the concentration directly:

Around 2%, but some insurers will say up to 5%, of claims by volume equate to about 85% of claims indemnity spent in the global property casualty market. The exposure is driven by low-frequency, but high-severity losses on individual risks.

The boundary-blindness condition imposes operational costs across the entire claims population, but these costs manifest differently at the two ends of the market. At the low-frequency, high-severity end (multi-policy losses, multi-jurisdictional placements, multi-party subscription markets, long-tail liabilities), the absence of portable decision evidence translates most directly into concentrated indemnity exposure: a small share of claims where coordination demands are highest. At the high-frequency, low-severity end (personal lines, small commercial, mass-market health and travel), the same gap surfaces at scale through AI-mediated decisions, outsourced TPA arrangements, and Consumer Duty's per-decision evidence requirements, with costs aggregated as regulatory penalties, complaint volumes, audit findings, and per-decision defensibility burdens. Both registers, concentrated indemnity at one end and aggregated regulatory burden at the other, translate into financial consequence; they do so through different mechanisms.

The condition appears at three boundary types that compound. Cross-firm boundaries (lead-to-follower, managing agent-to-coverholder, insurer-to-third-party administrator, primary-to-reinsurer) are the most familiar. Cross-platform boundaries (between distinct software systems, often within a single firm) produce structurally identical failures. Cross-jurisdictional boundaries (between distinct regulatory regimes, each with its own evidence-acceptance criteria) extend the condition to international reinsurance, multi-jurisdictional MGAs, and cross-border claims.

A fourth boundary, the insurer-to-customer boundary, sits adjacent to the three named above and is not treated in this paper. The same architectural condition appears at that boundary in the form of Consumer Duty, the Which? Super-Complaint and the Financial Conduct Authority's response⁸, but the present research was scoped to cross-firm coordination among regulated participants. The customer-facing dimension is treated in the underlying thesis, *The Aetiology of Boundary Blindness*.

A claims handler at one insurer working across three platforms (new-business underwriting, renewal underwriting, claims) faces the same boundary-blindness condition a follower syndicate experiences when the lead has performed decision work on a separate system. The dimensions compound: where evidence sits in Platform X and Platform Y, it cannot travel cross-firm even when both firms intend it to; and where those firms sit in different jurisdictions, the regulatory framework for that evidence is itself in dispute.

⁸ Which? (2025) Super-Complaint...The super-complaint triggered the FCA's December 2025 response and the subsequent expansion of FCA delegated authority oversight in Q1 2026.

A senior research participant accepted the term "boundary blindness" as a useful operational framework in April 2026, and it is now in active use across the research conversations. The remainder of this paper sets out the nine symptoms where the condition produces measurable cost in the present, the twenty-one published sources where it is being independently named, and the provisional architectural implication that follows.

2. Group A — Cross-Firm Coordination Symptoms

Five symptoms, all manifestations of the same structural condition: *decision context* fails to travel across firm boundaries, and across the platform boundaries accompanying them, in a form the receiving party can rely upon. The regulatory forcing functions that make these symptoms operationally intolerable are treated in Section 7.

S1. Followers cannot see what the lead reviewed

The lead-follow pattern appears wherever one regulated insurance participant relies on another's decision work: Lloyd's subscription (lead syndicate plus followers), broker-to-MGA-to-coverholder delegation, primary-to-reinsurer placement, and insurer-to-TPA claims arrangements. In each case, the follower commits capital, exposure, or a service obligation on the same terms as the lead but conducts minimal independent verification. The arrangement was efficient when relationships were collocated, and trust was bilateral. It has not aged well. And the same gap appears at scale in volume insurance, where AI-mediated claims decisions are passed between platforms, TPAs, and outsourced handlers under the same trust-the-prior-decision logic.

A Claims Director at a specialist advisory firm, working across multiple syndicates and complex commercial losses, described the operational reality directly:

You are the leader, you want to control the narrative. You don't want everyone to make comments before you make the comments.

Followers either accept the lead's decision regarding relational trust (which is increasingly difficult under regulatory scrutiny) or duplicate the work, often poorly, because primary evidence records are inaccessible. A claims handler at an independent UK regional brokerage described the same condition at the practitioner level:

As soon as another representative is added to a claim file, it is almost like starting from scratch; they are expected to gather the information.

A claims platforms consultant with three decades of London Market experience named the historical pivot:

You can't just, as I learned years before, [have] all the followers just try to trust the lead. You can't do that anymore. The days of the handshake — and we'll work it out — are gone.

The cross-firm coordination problem scales with the number of participants on a single risk. Humphrey, CILA Immediate Past President, described a current live claim from the loss-adjusting vantage:

On one claim, we have 55 participating insurers. A significant proportion of our time is therefore spent coordinating and aligning these stakeholders, ensuring all parties progress at the same pace and remain engaged throughout the process. The objective is to avoid fragmented decision-making and ensure no stakeholder feels excluded from decisions at critical points.

Humphrey identified a deeper structural challenge: the slip can mask the true allocation of economic risk. He referenced a recent major loss where the second-largest insurer on the slip carried a substantial 20% line. When questioned about their role in the decision-making process, their position was:

'Of our 20% line, we retrocede 19% to a global reinsurer.' That revealed an important dynamic: their interest was not solely in the underlying loss, but in ensuring the evidential basis and reporting were strong enough to sustain recovery further up the chain. In many respects, my role shifted towards protecting their reinsurance position.

The implication is structural. On paper, the slip-share is straightforward: each follower carries the percentage written next to their name, participates in decisions accordingly, and bears the loss in proportion to their share. In practice, retrocession often distorts that picture; a follower with a substantial line may retain only a small fraction of the economic exposure once their own reinsurance is taken into account. Their incentive then shifts: from reaching the most accurate loss reserve toward protecting the reinsurance recovery further up the chain. Other participants on the slip cannot see this. The lead-and-follower model presents one economic structure on its face and runs on another beneath the surface.

A complementary observation from the trade-body vantage extends the symptom. The Director of Claims at a London Market trade body distinguished two layers of the same gap: one, the follower's inability to verify the basis on which the lead reached a decision; the other, the follower's potential unawareness that a decision has been made at all. He characterised the second layer in worst-case terms:

I didn't even know you made a decision.

The two layers are distinct. The first is decision-quality awareness. The second is decision-existence awareness, a more fundamental gap that operates before the decision-quality question can be asked. The architectural answer must address both.

The same boundary-blindness condition is now visible in the AI vendor-deployer relationship. The Gallagher Re Q1 2026 InsurTech Report names the structural condition directly:

*"Beyond determining which policy responds, organisations must also consider whose policy responds. The legal and regulatory position of an organisation using AI often depends on whether it acts as an AI deployer... the AI deployer is generally the party responsible at law for the outputs of GenAI systems, with the AI treated as a tool rather than as an independent legal actor."*⁹

⁹ Gallagher Re, Gallagher, and CB Insights (2026), Global InsurTech Report 2026 Q1 — Artificial Intelligence: Risks and Opportunities in (Re)insurance and Beyond — Digital Risks v2, May 2026, pp. 62–63 (Risk allocation: The vendor-deployer divide).

The vendor-deployer divide is the cross-firm coordination gap in AI products: the vendor builds, the deployer uses, the deployer bears liability, and the vendor has the technical details. Decision evidence must travel from the vendor (build time) to the deployer (decision time) to the insurer (claim time)—the same architectural problem, expressed in AI-liability vocabulary.

The Lloyd's Principles Board Attestation Guidance for the 2026 supervisory cycle has elevated claims handling to a pass-or-fail attestation alongside underwriting and governance.¹⁰ Named individuals at managing-agent boards now carry personal accountability for demonstrating that lead-and-follower arrangements produce verifiable, defensible decisions across firm boundaries. Regulatory pressure has met an unaddressed structural condition.

Lead-follow remains legally compliant: the several-not-joint liability structure that defines the London Market subscription is unchanged. What has changed is the evidence regime that lead-follow must now satisfy. Followers historically relied on the lead under a relational trust model: the broker placed the risk, the lead syndicate underwrote it, the follower signed the slip and trusted that the lead would act competently on their behalf. That model has been quietly replaced since 2023, when the first material signals (the BoE/PRA SS1/23 model-risk principles, the NAIC Model Bulletin on AI Systems) began establishing the evidence regime that lead-follow must now satisfy. The Fifth Hurdle, the FCA's expanded delegated authority oversight programme, Solvency II exit plans, DORA, the EU AI Act, and APRA's Letter to Industry on AI each require something the relational model does not produce: documented, verifiable evidence that each participant, including each follower, exercised informed reliance on a documented standard at each material decision point.

The follower can no longer say "we trusted the lead"; the follower must say "we verified that the lead's decision met the documented standard, here is the record of that verification, here is the signed evidence." The signed acceptance signal is the artefact that makes informed reliance demonstrable. Without it, the follower's regulatory exposure scales with every decision the lead makes, because the follower cannot produce the evidence the new regime requires.

S2. Controlled asymmetry: leaders maintain narrative control

Cross-firm coordination is not failing through inattention. It functions because the asymmetry between lead and follower is largely deliberate.

The same Claims Director observed that leaders maintain narrative control because the alternative (every follower commenting on the lead's working notes in real time) produces operational chaos. She described a £6bn loss across Hull, P&I, pollution, liability, and energy, with up to fifty insurers across overlapping insurance towers, where uncoordinated follower commentary compounded the dispute rather than resolving it.

The same mechanism operates from the broker side. Humphrey, describing the day-to-day work of coordinating major-complex losses across the international placement chain, named it directly:

¹⁰ Lloyd's of London (2025) 2026 Principles Board Attestation Guidance. The framework elevates claims handling from a quantitative measure within the existing four-Principle structure to a fifth Principle requiring board-level attestation, colloquially.

In some cases, the brokers refuse to give you information about who sits behind whom. They may wish to be the conduit for that communication and not to facilitate communication between those parties. The broker's priority is to protect 'our client's interests'. And that can involve dividing and conquering.

The pattern is structurally identical to the lead-insurer pattern, at a different layer of the placement chain. The broker maintains it for a defensible reason: the client-fiduciary obligation. The mechanism is sustained by professional duty.

The same controlled-asymmetry mechanism, observed from the follower vantage, surfaces in the trade-body testimony:

I need to be comfortable with the decisions that the lead is making. If you're a claim follower, you can't just have that come through without understanding who it is that's making those decisions.

The follower accepts asymmetry as a structural necessity but requires confidence in the lead's decision quality. The architectural answer must therefore respect controlled asymmetry while producing a record the follower can use to verify the lead's decision-making against a documented standard.

The implication is direct. Transparency-by-default cannot be the design. The architectural answer must respect the controlled asymmetry the market has evolved for legitimate operational reasons: followers should verify that the lead has performed decision work to a documented standard, but in most cases, they should not read the lead's working notes. The same logic applies at the broker layer: participants in the chain should verify information has been handled to a documented standard, without each having visibility into every other party's relationship with the client.

The architecture preserves the lead-follow model that the subscription market has evolved for legitimate operational reasons. The lead retains the decision-making role; the change is that the follower can verify that the lead's decision has been carried out to a documented standard without reading the lead's working notes. The architecture does not displace lead-follow; it is made defensible under the new regulatory regime that demands evidence of informed reliance rather than blind trust.

S3. Selective disclosure is necessary

The Broking Director of an independent UK regional brokerage, supported by his Claims Handler colleague, raised the dashcam issue unprompted as a data protection concern. In a personal-lines motor claim, dashcam footage is itself relevant evidence (the motor incident is the matter being adjusted). Still, it frequently captures incidental personal data not relevant to the claim. The example given was the faces of dozens of schoolchildren on a passing bus. GDPR's data-minimisation principle prohibits incidental personal data from travelling through the chain alongside the relevant evidence; the brokerage has a positive obligation to ensure this does not occur.

Does that need to be shared with everyone? Probably not.

The Broking Director's concern is the commonest practitioner objection to cross-firm coordination infrastructure: not that it will fail to function, but that it will leak data and create regulatory exposure across overlapping regimes. The condition is operational reality across multiple regulated domains: the UK and European data-protection regimes¹¹ (GDPR and UK GDPR), the European Digital Operational Resilience Act,¹² the FATF Recommendation 16 Travel Rule for correspondent banking¹³ (cited here as a cross-firm-coordination architectural precedent rather than an insurance-applicable regime), and the confidentiality undertakings governing lead-follower relationships in the subscription market.

A second dimension of selective disclosure appears within the firm. A Data Protection Officer with over a decade in UK financial services data protection described the tension unprompted. The GDPR, retained in UK law alongside the Data Protection Act 2018, requires data to be held to the minimum necessary for a specific purpose. Financial services compliance and audit obligations require decision evidence be retained for years, sometimes decades, to defend against regulatory inspections or litigation. The two regimes pull in opposite directions: GDPR toward deletion, financial services regulation toward retention. The DPO put it directly:

He might want data retained for the next thirty years to enable us to defend any litigation. In my world, I do not want anything held, because as you are holding data, you are holding a risk.

A Compliance and Audit Manager with over four decades in UK financial services compliance, audit, and training, in the same joint interview, confirmed the trajectory:

Limitation-law test cases now pending could extend the standard six-year retention window to sixteen years, lengthening retention obligations at precisely the moment data-minimisation pressure is increasing. The two regimes are widening, not converging.

The Acceptance signal's separability property is the architectural reconciliation. The signed record of who reviewed the data, under what authority, with what evidence, and reaching what conclusion can be preserved over long retention horizons because it is structurally separable from the underlying data. The data can be minimised in line with its purpose-specific retention schedule, because the signed record provides verifiable proof of the decision chain without continued access to the source.

The architectural answer must therefore support disclosure granularity not only at the document level but within the document, at the field, region, or pixel level. Different recipients legitimately access different subsets of the same record. The infrastructure must be designed to respect existing regulatory regimes rather than retrofitted.

¹¹ Regulation (EU) 2016/679 (General Data Protection Regulation, "GDPR"), in force across the European Economic Area from 25 May 2018. Following the United Kingdom's withdrawal from the European Union, the GDPR was retained in UK domestic law as the "UK GDPR," read alongside the Data Protection Act 2018, with effect from 1 January 2021. Both regimes impose obligations on lawful basis, purpose limitation, data minimisation, and onward transfer that bear directly on the cross-firm portability of risk and claims information discussed in this paper.

¹² European Parliament and Council (2022) Regulation (EU) 2022/2554 on digital operational resilience for the financial sector (DORA). In force since January 2025 across European financial entities including insurance and reinsurance undertakings.

¹³ Financial Action Task Force (2023) Recommendation 16: Wire Transfers. The Travel Rule requires correspondent banks and virtual-asset service providers to share originator and beneficiary information across firm boundaries on qualifying transactions.

The market-infrastructure layer of insurance has independently arrived at the same selective-disclosure constraint and operationalised it in production. Lloyd's Delegated Data Manager (DDM), the service distributing risks, premiums, and claims data across managing agents, brokers, and follower markets today, implements participant-specific access as a built-in primitive of the data layer:

*"The Delegated Data Manager enables centrally controlled distribution of risks, premiums and claims data to participants, based on their allocated sections and markets in the binder, subject to the participant's access rights ... ensuring parties only see data which is relevant to them, adhering to data privacy standards."*¹⁴

The selective-disclosure constraint named here is not a researcher's overlay. It is the same constraint that the market-infrastructure community has deployed in production within a single market. This paper proposes the same primitive (role-aware, participant-aware, scope-aware access into a portable signed record) operating across the cross-firm, cross-platform, and cross-jurisdictional surface the nine symptoms occupy, rather than only within a single market's delegated-authority chain.

S4. Decision-state must hold at the moment of execution

Two convergence-cluster sources surfaced this symptom within ten days of each other in early May 2026: an independent advisor to property-and-casualty insurance executives, in a public translation of a manufacturing-AI essay; and a senior insurance commentator citing the same essay alongside the April 2026 APRA Letter to Industry on AI.¹⁵

The advisor's distillation was the cleanest plain-English version of the load-bearing claim of this research:

"The signal is not the decision. The decision is only real when it holds under consequence."

Decisions in insurance (binding coverage, paying or denying a claim, moving a reserve) are routinely contested months or years later. At execution, the decision must include the elements that enable its defence: context, authority, evidence, and conclusion. Where those elements are not captured at execution, the decision can be reconstructed from memory but not defended from the record.

Eric Daimler, former Senior Advisor for Artificial Intelligence to the Obama White House National Security Council and now Chief Executive of Conexus AI, makes a parallel observation about AI deployments in a public April 2026 essay:

*"When a federal regime retreats, liability does not vanish. It relocates. The question for any deploying institution is which direction it relocates."*¹⁶

The same record governs the relocation. Where decision evidence exists at the time of execution, liability falls on the party whose attested decision the record identifies; where it does not, liability falls on whichever party remains reachable later, typically the firm that deployed the model, not the firm that produced it.

¹⁴ Lloyd's of London. (2026). Delegated Data Manager. Market Resources / Tools. Retrieved from <https://www.lloyds.com/market-resources/tools/delegated-data-manager>(May 2026).

¹⁵ <https://www.apra.gov.au/apra-letter-to-industry-on-artificial-intelligence-ai>

¹⁶ Daimler, Eric (2026) "The American AI Accountability Gap." Public essay, Substack, 30 April 2026.

The same architectural deficit, expressed in AI-policy vocabulary.

S5. Reconstruction problem when decisions are contested years later

Where S4 concerns the moment of execution, S5 concerns the survival of decision evidence through time. In long-tail lines (Employers' Liability, Public Liability, D&O, professional indemnity, reinsurance casualty treaty business), claims emerge years or decades after the underwriting decision. When reconstructed for defence, audit, reserve review, or commutation, the original reasoning is often unrecoverable.

A senior commercial-lines underwriter with three decades of executive experience and a seat on the CII Underwriting Community Board named the deficit from the underwriting vantage:

An excess insurer would have to keep [a] permanent [file], including their own quote, the lead quote — and if you're an excess insurer, the underlying quote — because otherwise, you don't know what the market is doing beneath you.

The information that should travel between underwriting layers does not.

A senior loss adjuster with three decades in commercial property and specialty claims made the same point from the claims side. The original reasoning, she observed, typically lives only in the practitioner's head; defending the decision years later becomes especially difficult when supporting communications are missing from the file:

A lot of the reasoning is in their head... [where] there's not documentation to substantiate the thought process — if there's emails or phone notes missing in a file — that often makes it harder to review a file.

The same structural deficit reaches its most acute form at the long-tail edge, in reinsurance runoff portfolios where decisions made decades ago must be reconstructed for commutation, dispute, or settlement. A senior reinsurance business consultant with twenty-seven years in the London reinsurance market, twenty-three of them adjusting runoff and commutations, observed that the reinsurance market has historically been resistant to handing structured data forward through the chain:

Historically, people have never wanted to — or have been very, super reluctant to — hand over [data and decision authority] to the broker. They want to see everything come in and enter it themselves... [It is] an industry which is perhaps, certainly in the past, [been] guilty of being a bit of a dinosaur.

The reinsurance vantage extends the diagnosis to the long-tail end: the boundary-blindness condition is not only an operational friction but also the unresolved liability of three decades of unstructured decision records within firm-internal systems, now reaching maturity in runoff portfolios where they will be litigated, commuted, or written off without ever having been defensibly recorded.

The cross-vertical reinforcement is unusually strong. Munich Re's aiSure whitepaper¹⁷ documents the parallel failure in the AI supply chain, in which decisions made by AI systems will be revisited in litigation years later. The Geneva Association's analysis of fourteen public-private insurance programmes¹⁸ names the same failure at the public-private boundary.

3. Group B — Within-Firm Coordination Symptoms

Two symptoms in which the same structural condition appears within a single firm, frequently across the software platform boundaries that organise the firm's internal workflows.

S6. Cross-function friction inside a single firm

The structural condition the paper has described across firms also appears inside the firm, on a smaller scale, between functions that do not share decision evidence with one another.

A senior commercial-lines underwriter with over three decades of executive experience and a seat on the Chartered Insurance Institute Underwriting Community Board observed that the market uses the phrase "*separation of church and state*" to describe a governance principle: claims decisions are properly made by claims handlers, with underwriters consulted only when policy wording is ambiguous or a claims decision is challenged. The principle is normative and not in dispute. The operational reality is that the systems supporting underwriting and claims work vary by insurer; claims handlers may access underwriting files through processes that look different across firms. Whether the variable systems-separation makes the governance principle harder to operate in practice is the structural question the next practitioner voice in this section takes up.

The Broking Director of an independent United Kingdom regional brokerage, supported by his Claims Handler colleague, described the within-firm dimension in concrete terms:

The new business underwriter, the renewal underwriter, and the claims handler at the same insurer are working without a structured handover. Each starts almost from scratch.

The renewal underwriter does not inherit the decision work the new-business underwriter performed at placement. The claims handler does not inherit either. The same three-way breakdown (new-business underwriting, renewal underwriting, claims) surfaced in subsequent interviews with a Claims and Legal Services Director who chairs a community professional body, confirming that the pattern is not specific to any one insurer or segment.

The cost is borne by handler time, repeated information-gathering, and misalignments that lead to service complaints at the customer interface. Every coordination boundary, whether inside a single organisation or between distinct software systems, produces the same structural failure.

¹⁷ The five Munich Re Insure AI whitepapers: De-Risking Artificial-Intelligence Ventures (2023); Insuring Generative Artificial Intelligence (2024); Mind the Gap (2024, joint with Hartford Steam Boiler); From Governance to Insurance (2024); Generating Content with AI: An IP-infringement Minefield? (March 2026). Together

¹⁸ Geneva Association (2026). Addressing Growing Protection Gaps in Public-Private Insurance Programmes. Zurich: The Geneva Association.

A renewal underwriter able to verify the new-business underwriter's risk-acceptance reasoning, the rules applied, the evidence considered, and any claims activity since the prior renewal can price and write the renewal against the documented risk history rather than reconstructing it. The acceptance signal makes that decision history portable across the renewal boundary, even when the prior underwriter has left the firm. That is the workforce-transition condition S7 addresses at the structural level.

S7. Workforce transition and institutional memory loss

The London Market workforce currently numbers approximately 61,000 full-time equivalents and is ageing. The proportion of workers under 30 years of age was 24 per cent in 2024 and is projected to fall to 7 per cent by 2034 unless 4,400 additional graduates are hired per year, above the current rate.¹⁹ The London Market Group's "London Matters 2026 Fact Base" set out the trajectory in January 2026.

Comparable trajectories are visible in adjacent markets: US Bureau of Labour Statistics data show the US insurance workforce ageing along a similar curve, and Munich Re's 2026 Insure AI materials document similar demographic pressures across European reinsurers. The London Market figures are the most precisely measured; the structural condition they document is broader.

Boundary blindness has been operationally tolerable for three decades because experienced practitioners carried the institutional memory that filled the gap left by the lack of structured evidence. They could reconstruct the original decision-making reasoning when challenged because they remembered it or knew who to ask. The mechanism worked while two conditions held: enough experienced practitioners remained, and enough younger practitioners were entering to inherit what they knew. Both conditions are now breaking. Senior practitioners are ageing out, and the under-thirty cohort is itself collapsing from one quarter of the workforce to one-fourteenth. The absence of a structured evidence layer, tolerable for three decades, becomes acute at the workforce-transition layer.

The artificial-intelligence productivity proposition is the principal proposed offset to the workforce-transition gap. The London Market Group projects that AI productivity could reduce the graduate-hiring shortfall from 4,400 to approximately 1,800 per year in optimistic scenarios.²⁰ The condition for that offset is that AI systems can themselves operate on structured evidence inputs and outputs. That condition is not currently met. The same architectural deficit that produces boundary blindness in the cross-firm and cross-platform layers also disables the AI productivity layer intended to compensate for the workforce transition.

The next-generation testimony is consistent with this diagnosis. A senior London Market trade-body executive, describing his trade body's consultation with younger market entrants on the future-state market they want to inherit, observed:

¹⁹ London Market Group (2026) London Matters 2026 Fact Base, January 2026. The fact base is the principal industry source on London Market workforce demographics.

²⁰ London Market Group (2026), London Matters 2026 Fact Base, January 2026, workforce section. The productivity-offset scenario combines projected automation of routine task share with sustained graduate intake at current rates.

They [the broker community] still want to have a face-to-face relationship market — not fully anonymous, digitised, or [sit in front of] a computer — supported by electronic trading and electronic data-driven systems.

The same posture is voiced directly by the incoming cohort. A graduate-scheme market entrant in their early twenties, currently rotating through underwriting, claims, and distribution at a regional and SME insurer, described the appeal of the industry — and the design constraint that appeal places on any technology-led transition:²¹

Everyone always says that insurance is a people industry, built on relationships — whether that's right or wrong, that's how it is, and that's the biggest selling point of joining it. If everything sat within one system and you didn't need to talk to anyone, half the people employed would be out. What is the point of employing everyone if it is no longer going to be a relationship-based business?

The framing matters. The incoming generation is not asking for a fully digitised market that replaces human relationships. Nor is it asking for an unstructured market preserved by handshake. It is asking for something the industry has not yet produced: infrastructure that can carry decision context across the firm, platform, and jurisdictional boundaries those relationships span.

The same architectural layer that would address boundary blindness is the layer that would let the relationship-based market survive under regulatory and AI pressure.

4. Group C — Artificial-Intelligence Symptoms

Two symptoms where the convergence-cluster signal is now strongest. The artificial-intelligence dimension of boundary blindness is not the central thesis of this research, but it is where the published literature has accelerated most visibly in 2026.

S8. Wrong-layer architecture: controls in policy, not in the system

The Australian Prudential Regulation Authority (APRA), in its 30 April 2026 Letter to Industry on Artificial Intelligence, named the structural condition with regulatory precision: *"In many cases, preventative controls were lacking, with entities relying primarily on policy direction or detective, after-the-fact measures, rather than enforceable technical restrictions or robust preventative controls."*²²

A Fractional CTO for specialty insurers and MGAs, in public LinkedIn commentary on the APRA letter (6 May 2026), restated the argument from the practitioner vantage:

²¹ Primary-research interview, 29 April 2026. One of 61 senior-practitioner interviews conducted between September 2025 and May 2026 for this research. Notes on file with the author. Interviewee anonymised at source; identifying details (firm, exact role) omitted to preserve confidentiality.

²² Australian Prudential Regulation Authority (2026), Letter to Industry on Artificial Intelligence (AI) — Summary of common weaknesses and expectations for regulated entities. Sydney: APRA, 30 April 2026.

*"APRA is not saying governance is absent. It is saying governance is sitting at the wrong layer, primarily in documents and after-the-fact detection rather than in enforceable technical restrictions built into the execution path."*²³

A Compliance and Audit Manager with over four decades in UK financial services compliance, audit, and training restated APRA's findings from a UK practitioner's vantage point. Asked which is harder to close, policy-level controls or operational evidence, the response was direct:

Unfortunately, it is the latter. The policy is in place. They are in place in the business. But at the end of the day, it is up to how we actually evidence it and how we can justify that policy is actually working within the workplace. That is probably the one which is the hardest to do.

A forty-year UK compliance-and-audit vantage and an Australian prudential-regulator vantage reach the same diagnosis through independent routes: the policy-instrument layer is in place; the operational-evidence layer, where the policy is demonstrated to have been applied at decision time, is the harder gap where the missing infrastructure would sit.

Policy is not control. A decision supported only by written policy can be challenged on the basis that the policy was not applied at execution. A decision supported by a signed, structured record (produced at the moment of decision) is defensible by reference to the record itself. The architectural implication is that where controls are encoded matters more than what the controls say.

S9. Artificial-intelligence decision provenance gap

The convergence-cluster signal on this symptom is unusually dense. Sources across regulatory, AI-policy, reinsurer, and insurance-technology vendor sectors independently name variants of the same gap: a structural artefact recording what an AI system did, what a human did with that output, under what authority, with what evidence, and with what conclusion does not yet exist in operational form across firm boundaries.

Algorithmic decision-making is not new to insurance; it has been routine in personal-lines underwriting (home, motor) and in adjacent consumer-credit lending (mortgages, personal loans) for decades. What is new is its arrival in commercial and specialty markets, where decisions are higher-value, more contested, and must be defensible under regulatory and contractual standards that personal-lines decisions are not subject to in the same form.

APRA named the problem directly in its 30 April 2026 letter: *"Identity and access management capabilities have not yet adjusted to nonhuman actors such as AI agents."*²⁴

The Centre for Security and Emerging Technology, in a May 2026 issue brief,²⁵ reframed the problem mathematically. Within a Dempster-Shafer evidence framework, Belief and Plausibility bracket what is

²³ LinkedIn post dated 6 May 2026 commenting on the APRA Letter to Industry on Artificial Intelligence. URL:

<https://www.apra.gov.au/news-and-publications/apra-calls-for-a-step-change-ai-related-risk-management-and-governance>. The author serves as a Fractional Chief Technology Officer to specialty insurers and managing general agents; the role descriptor is reproduced at the level of generality at which the post itself is publicly available.

²⁴ APRA (2026), "AI threats are increasing, but information security capabilities are struggling to keep pace."

²⁵ Lohn, Andrew J. (2026) Beyond P(doom) for AI Risk: Quantifying Uncertainty Without Probability. CSET, May 2026. The formal claim cited is articulated within a Dempster-Shafer evidence-theory framework.

known about a proposition — the gap between them measures ignorance. Lohn writes: "Without ignorance, Belief and Plausibility become the same number, equal to probability." Read in the other direction: where epistemic uncertainty (ignorance) dominates aleatoric uncertainty (randomness), the two diverge — and probability, which collapses them into one number, becomes the wrong tool.

A practitioner-CEO in insurance data infrastructure, in a public post dated 1 May 2026, named the mechanism by which AI accountability evidence vanishes when the medium changes:

*"The loop is where accountability lives. It is also where mistakes get caught before they become events. Voice AI removes that loop. By the time someone could step in, it has already happened. That is not a flaw. It is the medium."*²⁶

The structural artefact required to evidence such decisions is not a slow review-stage record but a moment-of-utterance capture.

A Head of Claims at a major UK personal-lines insurer, with three decades in the industry and sat on industry bodies, framed the same problem from the practitioner vantage:

Artificial intelligence is something that is being used, but with caution. Being able to evidence fair outcomes for customers when you use data from artificial intelligence is on most claims leaders' minds at the moment.

The provenance requirement applies even where AI is semi-automated rather than fully agentic. A human reviewer who acts on AI-generated recommendations without recording the context available, the evidence reviewed, and the authority under which the decision was made carries the same defensibility gap as an AI agent operating alone.

Brian Perlstein and Michael Carroll, in a cross-vertical analysis from manufacturing operations published 5 May 2026,²⁷ named the failure mode plainly: *"A human can be in the loop and still not be judging. If the person approving the AI recommendation lacks the context, time, authority, and evidence to challenge it, we have not preserved judgment. We have preserved liability."*

The architectural artefact required to evidence judgment rather than approval is a record of the context available, the authority held, the evidence reviewed, and the conclusion reached at the moment of human disposition.

A final dimension of the AI decision-provenance gap concerns the temporal stability of the AI decision artefact itself. The same Group Data Protection Officer was asked what the evidence layer would need to support an AI-augmented decision that is defensible under both the UK GDPR and the EU AI Act. The response identified a dimension prior literature has not addressed: the model is non-stationary, and a decision recorded today against a given model state is not the same artefact as one the model would produce against the same input a week later:

²⁶ Public LinkedIn post dated 1 May 2026 by the Chief Executive of a company providing multi-party insurance data infrastructure. The post is cited at the speaker's level of generality consistent with the anonymisation discipline of this paper. Date and platform are recorded for evidential reference; full attribution is held by the author and available to qualified parties on request. URL on file.

²⁷ Perlstein, Brian, and Michael Carroll (2026) The Signal Is Not the Decision. Long-form essay, 5 May 2026.

If I put that same question in now, the answer that I am going to get, am I going to get the same answer next week? Based on artificial intelligence learning. The rates at which artificial intelligence learns could be by the minute, by the day, by the hour, by the week.

The architectural implication is that the Acceptance Signal must capture not only the AI output, the human disposition, the rules applied, and the evidence considered, but also the model state at the moment of decision: version, training cutoff, and configuration parameters. Without that capture, an audit months later cannot reconstruct the artefact that produced the original output. The model has continued to learn; the artefact of record is no longer instantiable.

Independent confirmation appeared in May 2026 in Gallagher Re's "Q1 2026 InsurTech Report." Ed Pocock, Global Head of Cybersecurity at Gallagher Re, named the same evaluation gap from a Tier-1 reinsurance broker vantage:

"Most AI models today are assessed using empirical benchmarks: standardised tests that measure accuracy on static datasets, performance on known tasks, or comparative rankings on public leaderboards. Recent releases have shown that models scoring well do not always behave consistently."

Pocock's analysis reaches the same conclusion from the underwriting-broker vantage point: a credible AI insurance market depends on an evaluation that focuses on behaviour, evolves, and reveals patterns of failure.²⁸ The acceptance signal architecture's requirement to capture model state at decision time is the evidence-layer expression of this requirement.

The provisional architectural implication: an evidence layer for AI decisions must be capturable across firm boundaries, must record both the AI output and the human disposition, and must survive transit through time at the level of regulatory defensibility.

5. The Convergence Cluster

Twenty-one independently published sources across six jurisdictions and seven sectors have named variants of the same architectural failure. Their publication dates cluster sharply (as of 10 May 2026): eleven of the twenty-one within the past three months, and seven of those eleven within the past three weeks. None references one another's work. Yet every one of them identifies the same point of failure: the boundary between firms.

The twenty-one sources, in order of publication date:

Sector	Source	Date	Jurisdiction
Insurance (Lloyd's)	Blueprint Two: Building the Future of Lloyd's	2020	UK

²⁸ Pocock, Ed (2026), Evaluating AI Risk: A Credible AI Insurance Market Depends on Better AI Model Evaluation, in Gallagher Re, Gallagher, and CB Insights, Global InsurTech Report 2026 Q1 — Artificial Intelligence: Risks and Opportunities in (Re)insurance and Beyond — Digital Risks v2, May 2026, pp. 42–48.

Sector	Source	Date	Jurisdiction
UK prudential regulation (model risk)	Bank of England / PRA Supervisory Statement 1/23, Model Risk Management Principles for Banks	17 May 2023 (in force May 2024)	UK
AI insurance underwriting	Munich Re Insure AI, five whitepapers	2023 to March 2026	Germany / global reinsurance
US insurance regulation	NAIC Model Bulletin on Use of Artificial Intelligence Systems by Insurers	December 2023 (model adopted; 2024–2026)	US
EU prudential regulation (insurance)	EIOPA, Opinion on Artificial Intelligence Governance and Risk Management in Insurance	2024	EU
US digital public infrastructure policy	Carole House (Atlantic Council Geoeconomics Center), Designing a Blueprint for Open, Free and Trustworthy Digital Economies	2024	US
London Market claims technology	DOCOSoft, Rising to the Challenge of Complex Claims	2024	UK / London Market
UK loss-adjusting trade body	Humphrey, Chartered Institute of Loss Adjusters President's Speech (CILA Annual Lunch)	11 April 2025	UK
London Market broker side	LIIBA, The Innovation Imperative	October 2025	UK
P&C insurance core platform	Guidewire, 2026 P&C Insurance Trends: The Forces Reshaping the Industry	22 December 2025	US / Global
Public-private insurance	Geneva Association (Schernberg), Addressing Growing Protection Gaps	February 2026	Switzerland / Global
Insurance industry standards (M&A)	ACORD, Carrier Mergers & Acquisitions: Drivers, Implications & Outcomes	2026	US / Global
Insurance industry standards (digital maturity)	ACORD, Insurance Digital Maturity Study: Creating Sustainable Advantage in a Digital World, 2026 Edition	April 2026	US / Global
Reinsurance trend forecasting	Munich Re Group + ERGO, Tech Trend Radar 2026	15 April 2026	Germany
Enterprise AI deployment	Cory Smith / Signal & Horizon, The AI Fever Dream	26 April 2026	US
Bank treasury	Eurobase International + Siena Treasury IQ, Treasury Resilience Checklist v1.0	27 April 2026	UK
Australian prudential regulation	APRA, Letter to Industry on AI	30 April 2026	Australia
Insurance data infrastructure	Practitioner-CEO public post (anonymised)	1 May 2026	US
Academic AI-risk policy	Lohn (CSET), Beyond P(doom) for AI Risk	1 May 2026	US
Manufacturing and industrial AI	Perlstein and Carroll, The Signal Is Not the Decision	5 May 2026	US
Global reinsurance broking (AI liability)	Gallagher Re, Gallagher, and CB Insights, Global InsurTech Report 2026 Q1: AI Risks and Opportunities — Digital Risks	May 2026	UK / Global

Each source identifies the family of symptoms documented in this paper. None names the cross-firm condition that produces them. The cluster spans 2020 to May 2026, with the data infrastructure and model governance preconditions established earliest (Lloyd's Blueprint Two; BoE/PRA SS 1/23; NAIC; EIOPA) and the practitioner experience, AI policy, reinsurer, and platform vendor voices concentrating sharply in the eight months immediately preceding this paper. The temporal density of the recent cluster is itself a structural finding.

The regulatory implementation of the cluster's institutional signals has accelerated in parallel. As of March 2026, the NAIC Model Bulletin had been adopted by twenty-five United States state insurance departments. Twelve states are concurrently running the NAIC Artificial Intelligence Systems Evaluation Tool pilot from March through September 2026, a structured framework that state examiners now use to review insurers' AI governance during market conduct examinations.²⁹

Mariusz Zagajewski, advising United States property and casualty Chief Underwriting Officers on Bulletin compliance, observed the operational gap directly:

"The breaking point of traditional underwriting is that this documentation is reconstructed under deadline pressure during an examination. An underwriting workbench produces it as a byproduct of operations."

Zagajewski's framing — documentation as a byproduct of operations rather than a retrospective deliverable — is the acceptance-signal pattern in underwriting-workbench vocabulary. The NAIC implementation creates the demand surface; the architectural answer it requires is the evidence layer this paper specifies.

A further dimension of public recognition has surfaced since the cluster was catalogued. In May 2026, three senior practitioners — Khalid Desai, Chief Data Officer at Allianz Global Corporate & Specialty (AGCS); Oliver Parker, Chief Technology Officer for Financial Services at DataArt; and Arno de Wever, Head of Commercial P&C Insurance UK at Amazon Web Services (AWS) — publicly diagnosed the within-firm fragmentation problem at an industry event.³⁰

They named 100+ source systems in commercial insurance, 23 policy administration systems per typical commercial insurer, and the failure mode of artificial-intelligence deployments built on unresolved data foundations. The medallion architecture (bronze, silver, and gold layers) was named as the within-firm answer. The cross-firm dimension was not addressed. The recognition has therefore reached senior data and technology communities serving the carrier population without yet reaching the structural diagnosis this paper names. The within-firm dimension is being diagnosed; the cross-firm and cross-jurisdiction dimensions remain unseen at the industry layer.

The convergence is mirrored by direct practitioner triangulation. A senior London Market trade-body executive, asked in a research conversation whether *decision rationale* is captured anywhere when

²⁹ Zagajewski, Mariusz (2026) "How an Underwriting Workbench Strengthens Regulatory Compliance: The 2026 NAIC AI Bulletin Playbook." Decerto, 9 May 2026 update of an August 2025 original.

³⁰ May 2026. The figures cited in the paragraph that follows — 100+ source systems typical of a commercial insurer, and 23 policy administration systems per typical commercial insurer — were presented by the three speakers during the panel discussion Data First: Building the Foundation AI Needs in Specialty Insurance (11 May 2026). All three are named in the body with their consent.

an underwriter binds a risk, confirmed that some capture exists — but locked inside firm-internal systems:

We have [encouraged the use of] new platforms and processes over the last few years — various, [the] Whitespace system or PPL — which actually has [an audit-able] step-through negotiation-and-bind process... that's where you kind of hand-type in how, why you made the decision, under what rules. The platforms record only the line, reference and details of the security — not the 'decision-making process.' Most of the [under]writers would do that in their own system rather than [the] central system.

The distinction the trade-body executive names — between the transactional data the platforms record (line, reference, security details) and the decision-making process they do not — is the architectural distinction this paper introduced in Section 1: data moves; decision context does not. The platform-vendor community has captured the data rails; the decision layer remains absent.

Asked whether the decision rationale then travels to the next platform or the next firm in the chain, he replied:

It doesn't, because each, other than the fact they've written it, suggests they've done it.

He continued, naming the structural reason no shared decision record exists:

Each line is several and not joint. The underwriting decisions are several and not joint.

The legal architecture of the London Market is the structural reason a portable, signed decision record has not historically been required. Each underwriter on a slip is liable only for their own line — "several" liability, not "joint." An underwriter with a 20% line is responsible for 20% of any loss, no more and no less. Because the law treats each underwriter's exposure as independent, no mechanism has ever required underwriters to share decision rationale with the other participants on the slip.

The same legal structure that supported subscription-market efficiency for three decades is now producing boundary blindness, as regulators, AI systems, and cross-jurisdictional supervisors require a shared record.

The London Market's four principal institutional voices have each published work in the ten months preceding this paper that names variants of the same diagnosis: the LIIBA Innovation Imperative (October 2025)³¹ on the broker side; the LMG London Matters 2026 Fact Base³² at the market-aggregate level; Lloyd's 2025 Annual Report on the managing-agent and Corporation side, stating that: *"Blueprint Two was a courageous undertaking...despite the efforts of many skilled and committed people, the project has not yielded the benefits that were originally envisioned. We have, therefore,*

³¹ LIIBA, Gracechurch Consulting and Trimstone Partners (2025) The Innovation Imperative: Why Brokers Matter More Than Ever, October 2025. Sponsored by Liberty Specialty Markets and Verisk. Page 46 calls for "shared schemas for claims data (akin to open banking APIs)."

³² London Market Group (2026) London Matters 2026 Fact Base, January 2026. Authors: Christopher Lay (Chair, LMG), Caroline Wagstaff (Chief Executive, LMG), and Graham Walsh (Campaigns Director, LMG).

taken the decision to sunset some of the project's original vision,"³³ and the International Underwriting Association (IUA)'s response to the HM Treasury Risk Transformation Regulations Consultation³⁴ on the company-market side. Each of the four corners of London Market governance is now publicly aligned on the diagnosis.

The IUA institutional position is corroborated by the practitioner-vantage interview for this research, confirming the trade body's diagnostic alignment from the company-market side. Each corner of the four-corner framing is now anchored at both the institutional/published and live-practitioner levels.

The broker-side institutional postmortem on Blueprint Two is the bluntest in the cluster. The same senior trade-body executive, asked what the broker community received from the programme, observed:

Brokers received [at] the bottom of this, just a very large bill and zero benefits. Those central services [that] Blueprint Two delivered will always [be] carrier-based services.

The asymmetric cost distribution is itself a structural finding. Any architectural answer must simultaneously produce participation incentives across the broker, carrier, bureau, and platform communities, or it will repeat the failure mode.

The platform-vendor side of the institutional triangulation was provided in April 2026 by a senior platform-technology executive with four decades of London Market experience across multiple major platform vendors. Asked what has actually changed in the underlying market data infrastructure over those four decades, he was direct:

Those electronic messages are a reimagining of [a record that] thirty years ago — forty years ago — was actually a punch card. The medium has changed, but the actual data is very poor.

The transport layer has now passed through four iterations of digitisation — from punch cards to electronic messages to APIs to AI-enabled extraction — while the structural problem it was supposed to solve has not changed. The observation extends the institutional convergence beyond regulators, trade bodies, and academic policy researchers into the platform-vendor community itself.

6. The Architectural Implication

The nine symptoms and the twenty-one cluster sources converge on a provisional architectural diagnosis: an evidence layer is structurally missing across regulated multi-party markets.

This paper proposes the term Portable Trust Infrastructure (PTI) to denote the category of infrastructure implied by the convergence cluster. PTI is infrastructure that sits alongside existing software platforms (claims systems, broker workflow systems, underwriting workbenches, document

³³ Lloyd's of London (2026), Annual Report 2025, page 9 (Chief Executive's statement on Blueprint Two). Corroborated on page 14 (Corporation strategic direction), page 79 (Council oversight of market infrastructure), and page 96 (March 2026 Blueprint Two reset announcement)."

³⁴ International Underwriting Association response to HM Treasury Changes to the Risk Transformation Regulations Consultation, 29 April 2026.

repositories, regulatory portals) and connects them with a new shared capability the existing software estate was not designed to provide: the production of portable, verifiable decision evidence.

PTI is therefore not a new platform, not centralised data aggregation, not an artificial-intelligence solution, and not a blockchain proposition. It is an event-driven coordination infrastructure that produces acceptance signals: structured, signed, portable decision records (defined in the Glossary) sealed by the cryptographic sigillum. These signals travel across firm, platform, and jurisdictional boundaries while preserving each firm's narrative control and selective-disclosure preferences, without touching the underlying data. The acceptance signal is the reference primitive of PTI; PTI is the architectural category.

Target Operating Model, Blueprint Two, B3i, and the Placing Platform Limited pre-Whitespace effort each required participants to agree on a shared data model and to share data through it. Each failed for the same political constraint: firms compete on data and resist harmonising shared models. PTI sidesteps this by design. It is not a shared data model but a verifiable record of decisions made against whatever data each firm holds. Firms retain their data, their schemas, and their commercial differentiation; what they share is the signed proof of decision, selectively disclosed. A PTI-conformant signal can travel through a harmonisation platform, alongside one, or in its absence, with the same verifiable properties in each case.

Independent articulation

The same architectural conclusion has surfaced from multiple vantages, in different vocabularies. Writing in the Insurance Innovation Reporter in April 2026, Rajesh Raheja named the diagnosis directly: *"That is not a technology gap. That is an architecture gap."*³⁵

He set out the answer in three clauses: *"Neural reads. Symbolic decides. Traces prove."*

He also distinguished a *"fluent rationale"* (the chain-of-thought output of an LLM) from a *"compliant decision trace"* (a record of how a decision was actually reached, anchored to evidence, rules, and authority). That is the same distinction this paper draws between AI output and the acceptance signal. The signal is the traces-prove layer; PTI is the category of infrastructure that produces it.

A year earlier, the same requirement had been named from the vantage point of the loss-adjusting trade body. Adam Humphrey, then-President of the Chartered Institute of Loss Adjusters, told the Institute's Annual Lunch in April 2025 that:

*"The aspiration remains to provide decision support to all stakeholders, entailing the delivery of high-quality information at the appropriate time to support good decision-making."*³⁶

³⁵ Raheja, Rajesh (2026) "Neuro-Symbolic AI Addresses Insurance's Explainability Challenge." Insurance Innovation Reporter, 28 April 2026. The piece independently arrives at the architectural diagnosis of this paper from a vendor-side vantage in the regulated-AI infrastructure space.

³⁶ Humphrey, Adam (2025) President's Speech, CILA Annual Lunch. London: Chartered Institute of Loss Adjusters, 11 April 2025. Speech delivered in the President's institutional capacity and shared with this paper's author 21 May 2026.

The three elements map onto the architectural specification: cross-firm portability, freshness and verifiability, and regulatory defensibility. These are the PTI properties expressed as practitioner aspiration a year before the category had a name.

Confirmation from the carrier vantage surfaced in May 2026. Matthew Busbee, Chief Data Officer of Pan-American Life Insurance Group, named the gap in operational vocabulary: *"the practical gap between 'we approved it' and 'we can continuously demonstrate it is behaving as intended,'" 37* and restated the vendor-deployer divide:

"Firms remain accountable for outcomes, even when the technology is delivered as a service."

The shift from a one-time approval at deployment to continuous demonstration over the life of the system is the same shift PTI specifies at the moment of decision and at every subsequent re-execution. Busbee is a third independent voice on the same accountability principle, alongside the Gallagher Re "Q1 2026 InsurTech Report" and the NAIC Model Bulletin.

The standards-body vantage

The published standards bodies have laid the data rails/tracks: ACORD³⁸ (insurance messages), the World Customs Organisation³⁹ (customs data), GS1⁴⁰ (product data), and ISO 20022⁴¹ (financial messaging). The data carriers exist: ECF, claims platforms, broker workflow systems, and partner-integration APIs. What does not yet exist is PTI.

A Senior Executive at ACORD, in a 2025 conversation with this researcher, named the missing layer in the standards body's own analogy:

ACORD is there to help with the actual laying of the data rails, but how you get in and out of that boxcar — that's a much bigger problem.⁴²

The same conversation named the cryptographic-verification design constraint:

³⁷ Busbee, Matthew (2026) "AI Governance: Creating A Responsible Framework." Pan-American Life Insurance Group, May 2026. The piece references the LIMRA and LOMA AI Governance Group (AIGG) seven-step Artificial Intelligence Project Lifecycle (AIPL) as the institutional scaffold for the framework Busbee describes. Pan-American Life is a United States life and accident insurer headquartered in New Orleans, Louisiana, with substantial Latin American and Caribbean market presence

³⁸ ACORD is the principal data-standards body for global insurance message formats.

³⁹ World Customs Organization, WCO Data Model 4.2.0 (current edition). Brussels: WCO. The WCO Data Model is the international standard for customs data exchange across border-control authorities and trade-facilitation parties, used by over 180 customs administrations globally. www.wcoomd.org

⁴⁰ GS1, GS1 General Specifications (current edition). Brussels: GS1 Global Office. GS1 maintains the global standards system for the unique identification of products, locations, parties and processes across supply chains, including the GTIN (Global Trade Item Number), GLN (Global Location Number) and SSCC (Serial Shipping Container Code) used internationally for product and shipment identification. www.gs1.org

⁴¹ International Organization for Standardization (2022), ISO 20022 — Universal financial industry message scheme. Geneva: ISO. ISO 20022 is the global standard for financial services messaging, used across payments, securities, foreign exchange, and trade-finance domains. www.iso20022.org

⁴² A Senior Executive at ACORD with public-facing responsibility for the development and dissemination of the insurance industry's data exchange standards. Quotation from a 2025 conversation with this researcher, reproduced at the speaker's level of generality, with sign-off on this descriptor from the speaker.

We want to allow creativity for what you stuff inside the boxcar. You or I don't need to dictate what it is — as long as it fits inside the boxcar, and somebody at the other end can trust that it wasn't tampered with.

In that vocabulary, PTI is the coordination layer that integrates four parts into a working whole. Two already exist: the data rails/tracks (the data-standards layer, such as ACORD) and the boxcars (the message and document containers). Two do not: the acceptance signal (the decision context sealed with the cryptographic *sigillum*) and the keys (the selective-disclosure access-control layer governing who can open, who can deposit, and who can read). PTI produces the missing two and integrates them with the existing two, operating as a coordination infrastructure rather than a centralised data infrastructure.

The contract-side counterpart was named by a senior London Market trade-body executive who has chaired oversight of the Market Reform Contract, the single artefact that, by design, crosses firm boundaries in the subscription market:

There has been a lot of people across the market [who] have tried to make the MRC the repository of [decision] data. It is not. It is not fit for purpose to do that. That is not what it's for... But it's just a framework for recording the details of a risk. It is nothing more [than] that.

The single contractual artefact that crosses firm boundaries today is not the carrier of decision rationale. PTI is the layer that the MRC was never designed to provide.

The practitioner side of the architectural answer surfaced independently. Humphrey, whose parallel work as an accredited mediator has placed him at the centre of distressed and disputed losses across the international placement chain, articulated three elements in operational vocabulary. On the architectural anchor:

Good practice is to establish a single source of truth so that all parties' decisions are based on the same facts. This supports an understanding of the points not in dispute, and therefore facilitates a focus on those that are and the reasons for the divergence of views.

And he named the selective-disclosure design constraint:

A secure data environment should be established for major losses, providing stakeholders with access to relevant information according to their roles and responsibilities. This creates transparency around decisions, establishes a clear and supportable audit trail, and enables technology to accelerate decision-making and deliver fairer, more equitable outcomes.

The three formulations describe PTI's defining properties (portable decision record, controlled asymmetry, selective disclosure) in practitioner vocabulary rather than category vocabulary. The diagnosis and the practitioner formulation are the same diagnosis, expressed in different vocabularies.

How the seal works in practice

The architecture above translates into recognisable operational scenarios. Three are illustrated below, each drawn from situations the practitioner cohort described as costly under current practice and catalogued in Sections 2–4.

How the seal is populated. Acceptance signals are generated automatically by existing platforms at the moment of each material decision. This includes when an underwriter binds, when a claims handler authorises payment, when a loss adjuster signs off on a report, and when an AI-mediated decision is reached. The platform's existing event flow (the same event flow that already writes to the CMS, the underwriting workbench, and the broker placement system) emits the signal as a byproduct of normal operations. Practitioners do not type anything into a new system or change their workflow. The information the signal carries (who, under what authority, applying what evidence, reaching what conclusion, at what point in time) is data the platform already has; the architecture proposed here captures it as a structured, signed, portable record rather than as scattered free-text in disparate logs. The integration work occurs once, between the platform and the acceptance-signal infrastructure; subsequent operation is invisible to end users.

Scenario A: External audit, three years after settlement. A managing agent receives a market conduct examination notice. The examiner selects a £4.2M commercial property claim settled two years prior. Under current practice, the audit response requires reconstructing the decision trail across the lead's claims management system, the loss adjuster's report archive, the broker's placement records, and the follower syndicates' own files, often through interviews with practitioners who have since retired or moved on. Under the seal architecture, the auditor receives a structured set of acceptance signals tied to the claim reference: the loss adjuster's signal recording delegated authority and the rules applied, the lead's coverage-decision signal, each follower's reliance signal demonstrating informed verification, and the settlement signal authorising payment. The audit response completes in days rather than weeks. Critically, the audit does not require the recovery of retired practitioners' recollections; the structured record speaks for itself.

Scenario B: A follower demonstrating informed reliance under FCA examination. A follower carries a 15% line on a complex commercial risk, with the lead having settled at £3.1M. The FCA's delegated-authority oversight programme selects the file for examination. Under current practice, the follower's ability to demonstrate "informed reliance" (as the Fifth Hurdle and the FCA's expanded oversight now require) depends on what their staff recall, what email correspondence survived, and whether they retained copies of the lead's working notes. Under the seal architecture, the follower's claims management system shows the signal-by-signal record: receipt of the lead's coverage-decision signal, verification against the lead's delegated-authority registry, review of the documented decision standard, and the follower's own countersignal authorising their share. The follower demonstrates informed reliance verifiably, without ever having needed to read the lead's underlying working notes. The shift from blind trust to verifiable reliance, as required by the new regulatory regime, is met structurally.

Scenario C: AI-mediated underwriting decision under EU AI Act examination. An insurer using an AI-assisted underwriting platform declines a high-risk submission. Under the EU AI Act Article 86 obligations (effective 2 August 2026), the applicant exercises their right to meaningful information

about the decision. The insurer must document not only that AI was used, but the specific basis on which this specific decision was reached. Under current practice, this requires reverse-engineering the model's behaviour from logs, model cards, and post hoc interpretability tools, none of which were available at the time of the decision. Under the seal architecture, the AI-mediated decision generated an acceptance signal at the time it was made: the model and version invoked, the inputs considered, the rules and thresholds applied, the score produced, the human reviewer (if any) and their override decision, and the conclusion reached. The audit trail EU AI Act, the NAIC AI Bulletin, and APRA's Letter to Industry each require that it be satisfied as a byproduct of operations, not reconstructed under deadline pressure.

Auditor and regulator access. Authorised auditors (internal compliance, external statutory auditors, and market conduct examiners) receive access to acceptance signals through their existing access pathway to the host platform, extended to expose the signal layer. The access model is permission-tiered: an internal auditor sees signals from their own organisation's activity; an external statutory auditor sees signals relevant to the audit scope under existing audit-rights provisions; a market conduct examiner sees signals tied to the examined file or portfolio under existing supervisory authority. The architecture does not introduce new access regimes; it makes existing audit and supervisory access more efficient by replacing scattered reconstruction with structured, signed, time-stamped records.

The three scenarios above illustrate the seal's operational behaviour at the firm level. The same primitive operates across the three boundary types unpacked below: a follower in a different firm verifies the lead's decision (cross-firm); a renewal underwriter in the same firm verifies the new-business underwriter's decision recorded on a separate platform (cross-platform); a European reinsurer verifies a US cedent's decision recorded under a different regulatory regime (cross-jurisdictional). The fields the signal carries are the same in each case; what changes is the receiving party and the verification context.

Boundary types and required properties

PTI must address three boundary types simultaneously: cross-firm portability (the dimension this paper has emphasised), cross-platform portability (most directly felt by platforms), and cross-jurisdictional portability (made unavoidable by the convergence cluster spanning six jurisdictions).

On the cross-platform dimension, an executive with vantage across the three central London Market platforms (placing, broker systems, bureau accounting and settlement) was direct:

The placing system is not fully connected to the broker system, which is not fully connected to [the] exchanging [accounting] systems. They don't loop, so there is not a true data transfer.

In the current market topology, the cross-jurisdictional dimension is encoded in who owns the platforms through which the data flows. The London Market's central bureau is US-domiciled: every premium transaction routed through it falls under US jurisdiction, regardless of where the underlying risk is written. UK and European underwriters who are legally entitled under their own national law to write Cuban risks cannot process those premiums through the bureau because US sanctions law applies

to every transaction involving a US-domiciled operator. The platform's jurisdiction overrides the jurisdiction of the decision. PTI must therefore be jurisdiction-neutral by design, so the acceptance signal travels independently of where the platforms it crosses are domiciled. The decision's jurisdictional posture must follow the decision, not the infrastructure it crosses.

Direct practitioner experience at the insurance/reinsurance interface reinforces the pattern. A senior loss adjuster on a current \$150M Middle East loss spanning three reinsurance layers described the operational consequence:

No one knows who to talk to. And so, the level of communication and the speed of decision-making are glacial because you're having to sort of run down a rabbit hole with information, but with no visibility.

The cross-jurisdictional dimension is not only a matter of differing regulatory regimes; it is a matter of who, within the placement chain, has visibility of whom, and on what terms. PTI must therefore expose not only the decision record but the decision-maker hierarchy, in a form that survives transit across jurisdictional layers.

The three boundary types are not separable in practice. A PTI-conformant signal that crosses firms but not platforms, or firms and platforms but not jurisdictions, leaves the problem partially solved. The specification must be platform-neutral and jurisdiction-neutral by design, supporting: cryptographic signing at the point of decision; field- and document-level selective disclosure; portability without centralisation; offline verifiability for auditors and regulators; survival of decision evidence over time; and machine-readability for AI systems reasoning from upstream decisions.

The bilateral-pilot risk

A counterpart pattern is already emerging in response to Blueprint Two's retirement. The same trade-body executive pointed to the most visible example:

Marsh and AXA XL have done a fantastic piece of work in conjunction with White Space, and have now delivered all the way through from placing right the way through to premium payments... using reusable and digitised data — which was the endgame of Blueprint Two.

The pattern matters in two directions. It confirms that Blueprint Two's architectural ambition was right: end-to-end portability of structured decision and data evidence across the placement-to-settlement chain. But it implies a structural risk: bilateral pilots solve the problem between the two firms involved without converging to a shared layer. As pilots multiply, the market accumulates point-to-point integrations rather than one shared infrastructure. The proliferation pattern is itself the next-generation manifestation of boundary blindness. PTI must be neither centralised (Blueprint Two's failure mode) nor bilateral (the emerging failure mode). It must be a shared, open specification that bilateral pilots can converge on.

Before any specification, two empirical questions need to be answered: whether the diagnosis correctly identifies what is missing, and whether a candidate PTI implementation can be tested in an operating environment with measurable results. Phase Four is structured to answer them.

What Phase Four will measure. The architecture described above produces value through three mechanisms that practitioner interviews surfaced consistently: reduction in audit-preparation effort (interviewees consistently described the current process as requiring weeks of practitioner time per major examination, with measurement of efficiency gains deferred to Phase Four); reduction in claim reopening and rework where decision context is lost in transit (interviewees described complex claims periodically requiring substantive rework attributable to coordination failures, with measurement of frequency deferred to Phase Four); and reduction in regulatory exposure for follower participants under the new evidence regime (quantitative impact dependent on the regulator and the examination outcome).

Phase Four will measure each of these mechanisms in the operating environments of three platform participants, with the findings forming the empirical basis for the architectural claim. Quantitative results are not asserted in this paper; the convergence is documented now, and the measurement is Phase Four.

7. Why Now: Three Forces Converging

Boundary blindness has been operationally tolerable for three decades because human institutional memory absorbed the cost of the missing structured-evidence layer. Practitioners reconstructed the missing decision context from what they remembered or from knowing who to ask. The mechanism worked while two conditions held: enough experienced practitioners remained to carry the memory, and the regulatory and technological environment moved slowly enough for human-paced workarounds to keep up. Three concurrent forces are now removing the absorber simultaneously: the workforce demographic transition is taking the memory itself; the regulatory cascade is demanding structured evidence that the memory was never able to produce; and artificial intelligence is scaling decisions at machine speed across boundaries the human absorber was never designed to cross.

Force one: the workforce demographic transition

The London Market workforce stands at approximately 61,000 full-time equivalents — and it is ageing. In 2024, 24% of those workers were under thirty; on current hiring rates, that share is projected to fall to 7% by 2034. Closing the gap would require an additional 4,400 graduates per year above the present intake. The over-fifties already account for more than a quarter of the workforce, and up to half of all practitioners are on track to retire within ten years. Senior practitioners are leaving the market faster than the under-thirty cohort can replace them. The institutional memory that has compensated for the missing structured-evidence layer is walking out the door.

The incoming cohort is not asking for the unstructured market to be preserved by handoff. A senior London Market trade-body executive, summarising consultations with younger market entrants, captured the ask directly: the relationship-based market *"supported by electronic trading and electronic data-driven systems"* — the structured layer the industry has not yet produced.

The same posture is voiced directly by the cohort itself. A graduate-scheme entrant in their early twenties, currently rotating through underwriting, claims, and distribution at a regional SME insurer, named the appeal of the industry and the design constraint that appeal places on any technology-led transition: *"Insurance is a people industry, built on relationships. If everything sat within one system,*

half the people employed would be out. What is the point if it is no longer a relationship-based business?"

The implication is direct. The incoming cohort cannot inherit the absorber's mode of compensation — informal, person-to-person, dependent on knowing who to call — because the absorber is leaving. What the cohort needs is infrastructure that produces the structured evidence as a byproduct of the relationship-based work they want to do, not a system that replaces it.

Force two: regulatory convergence

The second force is regulatory, and its distinguishing feature is compounding. Seven regulatory regimes — three of them coming into force on Lloyd's managing agents inside a six-month window in 2026 — each demand structured evidence that the human absorber was never producing. No single regime is the load-bearing one. Together, they require an evidence layer that no current Lloyd's market, European financial services, or United States prudential infrastructure produces.

The Fifth Hurdle — the Lloyd's 2026 supervisory cycle's pass-or-fail elevation of claims handling alongside underwriting and governance — came into force in January 2026.⁴³ The Financial Conduct Authority published its delegated authority oversight expansion through Q1 2026, following the Which? Super-Complaint of September 2025 and the FCA's resulting delegated-claims-handling enforcement programme.⁴⁴ Solvency II exit plans become required for Lloyd's managing agents in June 2026, including detailed claims strategies and evidence of cross-firm coordination under stress scenarios.⁴⁵

The Digital Operational Resilience Act has been in force across European financial entities since January 2025, with traceability, explainability, and human-oversight obligations for ICT-supported decisions in scope.⁴⁶ The EU Artificial Intelligence Act entered into force in August 2024 with explainability requirements for high-risk systems, including underwriting and claims decisioning.⁴⁷ APRA published its Letter to Industry on Artificial Intelligence on 30 April 2026.⁴⁸ The United States NAIC Model Bulletin on the Use of Artificial Intelligence Systems by Insurers had been adopted by twenty-five US state insurance departments by March 2026; twelve of those states are now running

⁴³ Lloyd's of London (2025) 2026 Principles Board Attestation Guidance. London: Lloyd's, 2025. The 2026 guidance elevated claims handling to a pass-or-fail attestation alongside underwriting and governance, with named-individual accountability at managing-agent boards.

⁴⁴ Financial Conduct Authority (2025) Response to the Which? Super-Complaint on Insurance Claims Handling. London: FCA, December 2025. The FCA's thematic review found that five of twenty-three firms examined had limited or no oversight of outsourced claims handling; the regulator initiated enforcement investigations, ordered skilled person reviews under Section 166 of FSMA, and expanded its delegated authority oversight programme through Q1 2026. The original Which? Super-Complaint of September 2025 (Which? 2025) triggered the review.

⁴⁵ HM Treasury (2026) Changes to the Risk Transformation Regulations 2017. London: HM Treasury, 2026. The amended regulations require Lloyd's managing agents to file Solvency II exit plans by June 2026; detailed claims strategies and evidence of cross-firm coordination under stress scenarios are a required component. See also Lloyd's of London (2024) Claims Lead Arrangements: Summary of Requirements, December 2024. coordination under stress scenarios are a required component. See also Lloyd's of London (2024) Claims Lead Arrangements: Summary of Requirements, December 2024.

⁴⁶ European Parliament and Council (2022) Regulation (EU) 2022/2554 on digital operational resilience for the financial sector (DORA). In force across European financial entities since January 2025. The regulation imposes traceability, explainability, and human-oversight obligations on ICT-supported decisions, including those produced by automated and AI-assisted systems used by financial entities and their third-party providers.

⁴⁷ European Parliament and Council (2024) Regulation (EU) 2024/1689 on artificial intelligence (AI Act). High-risk AI systems for which explainability requirements apply include underwriting, claims, and creditworthiness assessment. The traceability, explainability, and human-oversight obligations for high-risk systems take full effect from 2 August 2026; the structured evidence layer this paper names is the artefact those obligations require.

⁴⁸ Australian Prudential Regulation Authority (2026) Letter to Industry on Artificial Intelligence (AI) — Summary of common weaknesses and expectations for regulated entities. Sydney: APRA, 30 April 2026.

the NAIC Artificial Intelligence Systems Evaluation Tool pilot from March through September 2026, adding a seventh evidence regime at the US prudential layer.⁴⁹

Three of these regimes — the Fifth Hurdle (January 2026), the FCA delegated authority expansion (Q1 2026), and Solvency II exit plans (June 2026) — fall on Lloyd's managing agents within a six-month window. Three more (DORA, the EU AI Act, and APRA's Letter to Industry) compound the demand across European, EU AI explainability, and Australian prudential frameworks. The NAIC's pilot programme adds the US prudential layer. The convergence is not the appearance of a single new rule. It is the simultaneous demand from multiple jurisdictions, on overlapping populations of regulated entities, for evidence that the current infrastructure cannot produce, at a temporal density without precedent in the past three decades.

A claims platforms consultant with three decades of London Market experience named the qualitative shift that the convergence produces:

You're going to have to prove not only you did it, but that you did it with intent.

The shift the practitioner names — from action-evidence (that something happened) to intent-evidence (that someone intended it, applying what evidence, under what authority) — is the load-bearing change. Action-evidence the current infrastructure can already produce: a transaction log, a binder entry, and an ECF record. Intent-evidence requires the decision context that does not currently survive transit across firm boundaries. The new regimes do not allow one to be substituted for the other. Each requires, in its own regulatory vocabulary, the same structurally missing layer.

A senior executive at a London Market trade body highlighted the operational implications. Asked what happens when an external auditor — supervisory, regulatory, or third-party — needs to verify how a decision was made across the participants on a risk:

Once a contract has been entered into by two parties, we [the broker] sit in the middle. Someone would have to go to each individual underwriter and ask each one what their decision-making process was.

There is no architectural route by which an external auditor can reach the decision record without practitioner-by-practitioner reconstruction. The regulatory shift toward intent-based evidence collides directly with this. Where the auditor is the FCA, the EIOPA, the APRA, or a United States state insurance commissioner, the practitioner-by-practitioner reconstruction model fails at the moment of inspection — at the cadence the regulator now requires, across the boundaries the current architecture cannot bridge.

The compounding mechanism is structural. A single regulatory regime — even one as demanding as the Fifth Hurdle alone — could in principle be met by reorganising existing audit and quality-assurance functions within a single firm. The simultaneous imposition of seven regimes cannot. Each draws from

⁴⁹ National Association of Insurance Commissioners (2023) Model Bulletin on the Use of Artificial Intelligence Systems by Insurers. Kansas City: NAIC, December 2023. As of March 2026 the Model Bulletin had been adopted by twenty-five United States state insurance departments. Twelve of those states are concurrently running the NAIC Artificial Intelligence Systems Evaluation Tool pilot from March through September 2026 — a structured framework that state examiners use to review insurers' AI governance during market conduct examinations.

the same finite pool of management hours, audit capacity, and senior-practitioner attention; each requires evidence that the others' evidence frameworks do not satisfy; each operates on its own deadline. The market that has absorbed boundary blindness for three decades by routing the cost through institutional memory cannot route the cost through audit and compliance capacity at this volume and at this pace.

The same forcing function is visible on the platform vendor side. A senior platform-technology executive with four decades of cross-vendor London Market experience, hosting the heads of claims at industry roundtables in early 2026, described the buying-side shift:

The Fifth Hurdle has now been raised to the same level as underwriting and reserving. Suddenly, our customers who are heads of claims have got more ability to call upon the resources of their organisations, because it's a potential liability if they don't do this properly.

The regulatory shift has reorganised internal authority at the managing-agent and insurer levels. Heads of claims now command resources that were routinely directed elsewhere until 2026. The buying environment for the structural-evidence layer is therefore reshaped at the same moment that the demand for evidence is reshaping. The forcing functions create both the demand and the budget — a coincidence the past three decades did not produce.

Force three: artificial intelligence as a force multiplier

Artificial intelligence is scaling into claims and underwriting workflows at machine speed and across firm boundaries that human-paced workarounds were not designed to cross. AI was projected to offset the workforce shortfall — the London Market Group projects that AI productivity could reduce the graduate-hiring gap from 4,400 to approximately 1,800 per year in optimistic scenarios — but that offset is contingent on AI systems operating on structured-evidence inputs, a condition not currently met. The same architectural deficit that produces boundary blindness in the cross-firm and cross-platform layers also disables the AI productivity layer intended to compensate for the workforce transition.

Section 4 catalogued the institutional recognition of this dimension at the regulatory and academic layers. APRA, Munich Re Insure AI, the Centre for Security and Emerging Technology, and Perlstein and Carroll have each identified variants of the AI decision-provenance gap within their respective sectors. None has yet named the cross-firm condition that produces it. The artificial-intelligence accountability gap of 2026–2028 is boundary blindness scaled to machine speed.

The trade-body governance response in the UK has begun in parallel with the regulatory cascade. The Chartered Institute of Loss Adjusters' institutional restructuring under Royal Charter — Corporate Membership extending the Institute's Code of Conduct to insurers and brokers, and the name-change framing of claims as "*a team sport*" requiring "*collaboration of people in different organisations and multiple professional disciplines*" — represents a sector-side acknowledgement that decision-quality enforcement cannot remain firm-by-firm. Privy Council ratification of the restructuring is pending at the time of writing.⁵⁰

⁵⁰ Humphrey, Adam (2025) President's Speech, CILA Annual Lunch. London: Chartered Institute of Loss Adjusters, 11 April 2025. The speech identifies four CILA initiatives: Corporate Membership extending the Institute's Code of Conduct to a broader

The window

The cost absorbed over thirty years is about to surface in regulatory penalties, customer conduct complaints, audit findings, and competitive replacements. The window between absorption and surfacing is the window in which architectural intervention is still cheaper than remediation. The forcing-function analysis presented in this paper places that window at approximately 18 to 24 months from publication.

Two structural dynamics compound across that window. The first is network effects: each additional participant on a shared evidence layer increases the value of the layer for every other participant, so the platforms and firms that operationalise the acceptance signal early become the reference implementations against which later adoptions integrate. The second is regulatory positioning: participants who demonstrate the evidence layer ahead of the forcing functions cited in this paper establish a defensible position that later adopters must replicate under deadline pressure. The window is therefore not only the period when architectural intervention is cheaper than remediation — it is also the period when architectural choice is strategic rather than reactive.

8. What We Don't Yet Know

The research findings reported here are converging but not yet concluded. Several limitations are explicit, several empirical questions remain open to Phase Four and beyond; and several counterarguments require additional work before the architectural diagnosis can be considered settled.

Sample composition. Forty-seven of the sixty-one practitioner interviews were conducted in the insurance markets of the United Kingdom and the United States. Cross-jurisdictional practitioner depth is concentrated in those two regimes; Asian, Middle Eastern, and continental European practitioner voices are present at the institutional layer (Munich Re, Geneva Association, APRA) but thin at the operational-practitioner layer. The same imbalance applies to the personal-lines, community-broker, and direct-insurance segments: present in the sample but underweighted relative to Lloyd's and specialty representation. The diagnosis is robust in the practitioner segments where the sample is strongest; its generalisation to underweighted segments is provisional.

Architectural specification. This paper describes the missing layer—the *acceptance signal*—at both conceptual and functional levels: what it must do, the design constraints it must meet, and the boundary types it must cross. The paper does not provide a formal protocol specification, a security or threat model, a reference implementation, or interoperability test results against ACORD, ECF, or other existing market infrastructure. Phase Four is structured to test architectural fit against operating environments. It is not structured to test architectural correctness in formal terms. The full specification is a downstream artefact, not a finding of this paper.

S7 (workforce transition). Well-documented, but its trajectory remains uncertain. The London Market Group's projections depend heavily on artificial-intelligence productivity offsets that are themselves untested at scale. The 2034 trajectory the LMG sets out assumes the AI productivity offset materialises

stakeholder group; a new Honorary Membership category; institutional restructuring of governance approved by the Council; and the framing of claims as "a team sport" requiring cross-organisational collaboration. Privy Council approval is described in the speech as "underway, but...a time-consuming process"; the paper's "pending at the time of writing" reflects that status as of May 2026.

at scale. If that offset is materially smaller than projected, both the workforce-transition pressure and the demand for structured-evidence infrastructure arrive on a shorter timeline.

S9 (artificial-intelligence decision provenance). Theoretically pressing, but practitioner reports of AI in live claims decisions remain limited as of Q2-2026. A Claims Director at a specialist advisory firm distinguished between the digital-transformation work happening across the London Market in 2026 — process automation, document summarisation, claim-file reading — and AI used for decision-making, which she characterised as approximately a year away in an April 2026 conversation. Phase Four research will produce a clearer empirical signal.

Cross-industry generalisation. The cross-industry dimension is asserted but has been triangulated only at the institutional layer. Whether the architectural pattern that addresses boundary blindness in insurance generalises operationally to customs, logistics, banking, or healthcare remains a downstream research question. The operational mechanism by which an open specification would be adopted across vertical markets has not been examined here.

The research has, however, surfaced concrete examples of insurance itself depending on signals from adjacent industries. A FedEx shipment of temperature-sensitive medicine carries a sensor that records whether the consignment has exceeded its allowable temperature; an exceedance during the overseas flight can trigger a multi-million dollar insurance claim for a single shipment. The claim originates in logistics data — not in any insurance system — and the boundary problem documented in this paper appears at the handoff between the carrier, the cargo insurer, and the pharmaceutical shipper. The same research is also beginning to expose operational silos within these adjacent industries themselves: early evidence that the architectural pattern diagnosed here in insurance is not confined to insurance.

Coverage of the customer-facing boundary. The three-boundary diagnostic frame set out in Section 1 is deliberately scoped to cross-firm coordination between regulated participants. The insurer-to-customer boundary, mediated by the broker, is not treated in this paper. The interview sample includes practitioner voices from the personal-lines and community-broker segments, and incidental practitioner evidence surfaced during the research has been catalogued for the underlying thesis. Whether the boundary blindness condition operates with structurally identical symptoms on the insurer-to-customer boundary, and whether the acceptance signal architecture should treat the customer as a relying party with read access, a producer of consent attestations, or both, are downstream research questions for the thesis and for subsequent phases. They are not findings of this paper.

Counter-arguments requiring empirical work. The strongest counter-arguments to the architectural diagnosis are political-economic rather than technical. First: the controlled asymmetry this paper diagnoses as a structural feature is, for many participants, a structural benefit. Leads benefit from narrative control; brokers, from divide-and-conquer; lawyers, from disputes over interpretation. The Acceptance signal architecture preserves controlled asymmetry by design: selective disclosure means followers can verify that decision work was performed to a documented standard without reading the lead's working notes. But the shift from no portable verifiable record at all to a portable verifiable record subject to selective disclosure is itself a redistribution of information power. Participants who benefit today from the absence of any verifiable record, not from the leakage of specific content, may resist a

verifiable record regardless of how granular the disclosure controls are. The paper does not yet address which adoption-incentive structure would overcome these entrenched positions. Second: bilateral firm-to-firm pilots, with the Marsh / AXA XL / White Space arrangement as the most visible example, may converge organically on functional standards without requiring an open specification at the market layer. Whether bilateral proliferation produces de facto convergence, or whether the proliferation pattern itself becomes the next-generation manifestation of boundary blindness, is empirically uncertain.

Buyer acceptability and the audit framing. A third counter-argument concerns whether practitioners would welcome a structured signed record or reject it as surveillance over their judgment. Earlier interview testimony in this research surfaced the surveillance objection: at least one practitioner described the framing as a challenge to judgment rather than a coordination tool.

The objection has now been examined from the trade body's institutional vantage point. The Director of Claims at a London Market trade body asked directly whether members would view a structured, signed record as a useful coordination tool or as surveillance, was direct:

I don't think that's something that members would necessarily view as an obstacle. I think they would view that as actually a quite useful tool.

The acceptance is anchored to the Insurance Act framework. The same source noted that members are already required to demonstrate how and why decisions were reached, particularly when claims are declined; the Insurance Act requires the insurer to show that the decision is one the insurer would have made under the circumstances. The structured, signed record, framed as Insurance Act audit evidence rather than as decision transparency that challenges judgment, lands as an enabling tool rather than an imposition. The architectural answer must produce audit-friendly evidence consistent with statutory and regulatory audit obligations, not surveillance-style transparency. The earlier surveillance objection has not vanished, but the practitioner population from which it would arise is materially narrower than the institutional acceptance suggests.

Adoption pathway and commercial framing. Which firm, which regulator, or which market layer would adopt the architectural pattern first is outside this paper's scope and is the question that will determine whether the diagnosis matters in practice. Pricing and commercial-value hypotheses for platform-mediated infrastructure are similarly not addressed. Phase Four research tests technical feasibility and architectural fit; any commercial implications emerge as findings of that technical work, not as prior assertions.

9. Phase Four — An Invitation

Phases one through three conducted senior-practitioner interviews, analysed convergence-cluster sources, and converged on the diagnostic and nine symptoms set out above. Phase Four is the empirical-validation step, conducted under the Verida Charter Foundation (a 501(c)(6) US nonprofit).

The Phase Four programme welcomes participation from multiple categories of organisations. Its empirical core engages three software platform participants in a structured research collaboration; the broader programme also invites participation from trade associations, standards bodies, research

foundations, government agencies, and individual practitioner experts. Each pathway has a defined contribution mode and reciprocal value.

The empirical core: three software platform participants

The empirical work engages three platform participants in a 12-week research collaboration designed to test two questions:

1. Whether the VCF-compatible Platform Connector Kit for each participating platform integrates with low friction in that platform's operating environment.
2. Whether the practitioner-validated symptoms catalogued in this paper resonate with the platform's own customer experience.

The structure of the platform engagement:

- Three platforms across the UK, European, and (where practical) US or Australian insurance domains, selected to represent different software environments, regulatory mixes, and customer bases.
- One full calendar quarter (approximately 12 weeks) per platform, with three-week staggered starts beginning July 2026: scoping (weeks 1–2); connector kit integration (weeks 3–4); workflow testing (weeks 5–7); edge-case refinement (weeks 8–9); cross-firm and cross-platform scenario testing (weeks 10–11); findings capture and exit (week 12).
- Approximately 25–30 s of platform engagement across the quarter, distributed across engineering, product team review, and the closing research conversation. Average 2 hours per week, scheduled at the participant's convenience.
- Platform Connector Kits are co-developed during the engagement to maintain compatibility with each platform's operational environment and the evolving VCF schemas. Each kit is proprietary to the platform that co-develops it and is not open-sourced; what is open is the VCF standard and the connector-conformance suite each kit is certified against. The kits are tested against the author's prototype reference implementation, the available VCF-compatible backend during the research period. The VCF API specification and the conformance suite are published under Apache 2.0 at the close of Phase Four, so that any party may build and certify an independent VCF-compatible backend.
- A cross-platform joint debrief midway through the engagement window, attended by all three platform partners (with permission), where each platform presents its independent findings. This is the moment when the universality claim is tested empirically: do three platforms, using different connector kits in different environments, yield consistent findings?
- No purchase or procurement commitment. Phase Four is research participation, not a commercial pilot. Completing Phase Four creates no obligation for the participating platform to engage commercially. Participating platforms retain full editorial control over the findings published about their participation, subject to anonymisation where requested.

What platform participants get:

- Outright ownership of the connector kit for their platform. The platform's engineering team co-develops the kit during Phase Four to maintain compatibility with the platform's operational environment and the evolving VCF schemas. The resulting kit is the platform's own proprietary asset, with full rights of use, maintenance, and future development in perpetuity; it is not released as open source. The Verida Charter Foundation does not take ownership of any platform's connector kit — the Foundation owns and publishes the open standard (the VCF Protocol, the Acceptance API specification, and the conformance suite) against which each kit is certified. The platform retains its kit regardless of any future decision regarding the prototype backend, including any transition to a commercial operator.
- Early visibility into the boundary-blindness diagnostic and the architectural response, with the opportunity to shape the connector kit specification for their platform.
- Direct research engagement with the author across the 25–30 hours of the quarterly engagement.
- An integration friction report for their platform; inclusion in the joint Problem-Fit Matrix; and the joint debrief recording.
- Named acknowledgement in the published research (subject to the platform's consent).

Other Phase Four participation pathways

The Phase Four programme also welcomes broader institutional participation from organisations whose role does not involve building or integrating platform-specific connector kits but whose engagement strengthens the research's institutional foundation, regulatory relevance, and standards convergence.

Trade associations (such as the Lloyd's Market Association, LIMOSS, LIIBA, the London Market Group, the Chartered Institute of Loss Adjusters, the International Underwriting Association, BITA, and equivalent bodies in adjacent jurisdictions) can participate as research sponsors, named endorsers, member-convening hosts, or contributors to the post-Phase Four standards governance of the VCF Protocol. Participants receive named acknowledgement, observer status at the cross-platform joint debrief (with permission), and an institutional voice in post-Phase Four governance.

Standards bodies (such as ACORD, GS1, the World Customs Organisation, and ISO) can participate through standards-alignment review, co-publication of conformance criteria, and structured engagement on the harmonisation surface between their existing standards and the VCF Protocol. Participants receive named acknowledgement and structured engagement on standards-alignment outputs.

Research foundations supporting open infrastructure, public-interest research, or regulated-market resilience can participate as Phase Four funders. Foundation funders receive named acknowledgement and access to research outputs on the same terms as any other party: Apache 2.0 for the open standard (the VCF Protocol, the Acceptance API specification, and the conformance

suite) and an open licence for the published findings. Connector kits are proprietary to the platforms that build them and are not part of the open release. (Apache 2.0 for code; open licence for findings).

Government agencies and regulators (such as the FCA, EIOPA, NAIC, APRA, NIST, NSF, DARPA, and EU Horizon) can participate as observer-funders, providing both research support and regulatory-context translation. Participation is subject to the standard research funding disclaimers: findings represent the author's analysis and do not necessarily reflect the views of any participating agency.

Individual practitioner experts whose expertise complements the 61-interview cohort can contribute through additional structured interviews, serving as reviewers on draft findings, or participating in post-Phase Four research collaboration. Participants receive named acknowledgement (subject to consent).

General expressions of interest. Organisations whose role does not fit cleanly into the categories above but whose engagement is relevant to the research are invited to make contact.

Deliverables and research outputs

Phase Four research outputs include: integration friction reports per platform; the joint Problem-Fit Matrix; the joint debrief recording (subject to platform consent); three or more platform-specific connector kits, each owned by its participating platform as a proprietary asset and not open-sourced; the VCF API specification and the connector-conformance suite, released under Apache 2.0; an empirical chapter in the underlying research thesis; and a public Phase Four findings paper targeted for Q1 2027, including any findings unfavourable to the architectural pattern or the author's prototype work.

Engagement framework

Phase Four operates under a structured research-engagement framework with safeguards including research-purpose-only access for platform participants, confidentiality protection over implementation details (architectural concepts in this paper are public and excluded from confidentiality scope), intellectual-property arrangements under which each platform owns its connector kit as a proprietary asset and the Verida Charter Foundation owns and publishes the open standard and conformance suite, cooling-off provisions, and an explicit no-obligation clause. For each participant category, the appropriate full participation agreement is provided upon submission of an indication of interest.

Specific invitations and indication of interest

Specific invitations are being issued in conjunction with this paper. If you have not received an invitation directly and your organisation has a role to play in the architectural questions this research addresses, please get in touch with the author. Indication of interest by 26 August 2026. Platform engagements complete between October and December 2026; broader institutional engagement is ongoing through the Phase Four period and beyond; findings published Q1 2027.

Phase Four — Empirical Validation

Three platform participants. Approximately 25 – 30 hours of engineering and product time per platform. Twelve-week sandbox per platform. Joint debrief in week seven. No commercial commitment. Indication of interest by 26 August 2026.

Alexander D. Barrett, Senior Fellow, Verida Charter Foundation — adb@veridafoundation.org

Methodology and Acknowledgements

This research was self-funded by the author and conducted between October 2025 and May 2026 under the auspices of the Verida Charter Foundation, a nonprofit research foundation focused on cross-firm coordination infrastructure in regulated multi-party markets. No external party funded or directed the work. The methodology comprises: sixty-one semi-structured interviews with senior practitioners (forty-seven within United Kingdom and United States insurance markets, fourteen across adjacent regulated multi-party domains), conducted under anonymisation undertaking; first-hand analysis of twenty-one published convergence-cluster sources across six jurisdictions and seven sectors; triangulation against contemporary trade-body publications; and parallel review of contemporary regulatory publications from the FCA, the PRA, APRA, and Lloyd's of London.

The research has also generated substantive engagement from senior practitioners independent of the original interviewee pool during the period of this revision. The Head of Legacy Origination at a Tier-1 global reinsurer initiated unsolicited inbound contact in response to the researcher's public engagement on an adjacent topic. Public exchanges with a Senior Fellow at a leading United States academic policy institution, the Chief Executive of a multi-party insurance data infrastructure company, and a graduate of a European business school now in international advisory practice extended the conversation into the United States, infrastructure-builder, and adjacent-domain communities. The pattern is consistent with the convergence set out in Section 5: the structural condition this paper diagnoses is being recognised at institutional and academic levels, in independent vocabularies, in the same window in which the paper has been prepared.

The research has also benefited from sustained cross-domain peer exchange. During the research window, the author convened with senior fellows and counterparts in adjacent regulated multi-party markets, sharing working notes across anti-money-laundering (AML) and financial-crime supervision, logistics and customs, global supply chains, capital markets and financial-services infrastructure, and the intelligence community. These exchanges tested whether the architectural diagnosis advanced in this paper, that decision evidence fails to travel across firm, platform, and jurisdictional boundaries, generalises beyond insurance. In each domain, senior counterparts independently recognised variants of the same condition within their own sectors, using their own vocabularies, without prompting from this researcher. The cross-domain triangulation does not displace the insurance-specific empirical work that anchors this paper; it strengthens the claim, advanced in Section 8, that the boundary-blindness condition is a structural feature of regulated multi-party markets generally, not of insurance specifically.

All practitioner quotations in this paper are anonymised at the role and tenure level. Public authors of published artefacts are named. The author holds full participant acknowledgements and is available to qualified parties on request, with each participant's explicit consent.

The author is profoundly grateful to the sixty-one practitioners — claims directors, underwriters, brokers, loss adjusters, platform leaders, trade-body executives, regulators-adjacent commentators, and senior voices from across the United Kingdom and the United States insurance, and from adjacent regulated multi-party domains — who gave their time, their candour, and the accumulated experience of careers spanning decades. Many sat for half an hour or longer during a working day; some returned for second and third conversations; several offered introductions that opened entirely new lines of enquiry. Senior figures within the Chartered Insurance Institute, the Lloyd's Market Association, the London and International Insurance Brokers' Association, the International Underwriting Association, and the Chartered Institute of Loss Adjusters made institutional access possible that no individual researcher could have reached alone.

The diagnostic this paper sets out, the nine symptoms it catalogues, the convergence pattern it identifies, and the architectural implication it draws are as much theirs as they are the author's. This paper is the synthesis of their observations; it could not exist without them. Any errors of synthesis are the author's alone.

The full methodology, regulatory analysis, and detailed evidence base are set out in the underlying research paper, "*The Aetiology of Boundary Blindness*," working thesis, May 2026, available on request.

Alexander Barrett | Verida Charter Foundation | Washington, D.C.
adb@veridafoundation.org | www.veridafoundation.org

References

1. ACORD (2026). Insurance Digital Maturity Study 2026 Edition. New York: ACORD.
2. Australian Prudential Regulation Authority (2026). Letter to Industry on Artificial Intelligence (AI) — Summary of common weaknesses and expectations for regulated entities. Sydney: APRA, 30 April 2026.
3. Bank of England (2023). Supervisory Statement 1/23: Model Risk Management Principles for Banks. London: Prudential Regulation Authority, 17 May 2023; implementation deadline 17 May 2024.
4. Busbee, Matthew (2026). "AI Governance: Creating A Responsible Framework." Pan-American Life Insurance Group, May 2026.
5. Daimler, Eric (2026). "The American AI Accountability Gap." Public essay, Substack, 30 April 2026. Available at: <https://ericdaimler.substack.com/p/the-american-ai-accountability-gap>
6. ERGO Group (2026). Tech Trend Radar 2026: Press Release, 15 April 2026.
7. Eurobase International and Siena Treasury IQ (2026). Treasury Resilience Checklist v1.0, 27 April 2026.

8. European Insurance and Occupational Pensions Authority (2024). Opinion on Artificial Intelligence Governance and Risk Management in Insurance. Frankfurt: EIOPA.
9. European Parliament and Council (2022). Regulation (EU) 2022/2554 on digital operational resilience for the financial sector (DORA). OJ L 333.
10. European Parliament and Council (2024). Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (AI Act). OJ L 2024/1689.
11. Financial Action Task Force (2023). Recommendation 16: Wire Transfers. Paris: FATF.
12. Financial Conduct Authority (2025) Response to the Which? Super-Complaint on Home and Travel Insurance Claims Handling. London: FCA.
13. Financial Conduct Authority (2026). Regulatory Priorities for Insurance 2026. London: FCA.
14. Gallagher Re, Gallagher, and CB Insights (2026). Global InsurTech Report 2026 Q1 — Artificial Intelligence: Risks and Opportunities in (Re)insurance and Beyond — Digital Risks v2. London / New York: Gallagher Re, May 2026. 112-page report; the last in a trilogy of AI-focused annual reports begun in 2024.
15. Geneva Association (2026). Addressing Growing Protection Gaps through Better Public-Private Insurance Programmes. Author: Hélène Schernberg. Box 7 contribution by David Bresch, ETH Zürich.
16. HM Treasury (2026). Changes to the Risk Transformation Regulations Consultation, 29 April 2026.
17. House, Carole (2024). "Designing a Blueprint for Open, Free and Trustworthy Digital Economies." Washington, DC: Atlantic Council Geoeconomics Centre.
18. LIIBA, Gracechurch Consulting and Trimstone Partners (2025). The Innovation Imperative: Why Brokers Matter More Than Ever. London: LIIBA, October 2025.
19. Lloyd's of London (2020). Blueprint Two: Building the Future of Lloyd's. London: Lloyd's.
20. Lloyd's of London (2024). Claims Lead Arrangements: Summary of Requirements. London: Lloyd's.
21. Lloyd's of London (2025). 2026 Principles Board Attestation Guidance. London: Lloyd's.
22. Lloyd's of London (2026). Annual Report 2025. London: Lloyd's, March 2026. See page 9 (Chief Executive's statement on Blueprint Two), page 14 (Corporation strategic direction), page 79 (Council oversight of market infrastructure), and page 96 (March 2026 Blueprint Two reset announcement).
23. Lohn, Andrew J. (2026). Beyond P(doom) for AI Risk: Quantifying Uncertainty Without Probability. CSET, May 2026.
24. London Market Group (2026) London Matters 2026 Fact Base. London: LMG.
25. Munich Re (2023–2026). Insure AI five-paper canon: De-Risking Artificial-Intelligence Ventures (2023); Insuring Generative Artificial Intelligence (2024); Mind the Gap (2024, with Hartford Steam Boiler); From Governance to Insurance (2024); Generating Content with AI: An IP-infringement Minefield? (March 2026).
26. Munich Re Group (2026) Tech Trend Radar 2026: Our expert perspective for insurance. Munich: Munich Re.
27. National Association of Insurance Commissioners (2023). Model Bulletin on Use of Artificial Intelligence Systems by Insurers. Kansas City, MO: NAIC, December 2023. State implementations 2024–2026.
28. Perlstein, Brian, and Michael Carroll (2026). The Signal Is Not the Decision. Long-form essay, 5 May 2026. (<https://www.linkedin.com/pulse/signal-decision-michael-carroll-lhc9e/?trackingId=5jD8MxfwC6OzJ46dWoyMJA%3D%3D>)
29. Raheja, Rajesh (2026). "Neuro-Symbolic AI Addresses Insurance's Explainability Challenge." Insurance Innovation Reporter, 28 April 2026.

30. Smith, Cory / Signal & Horizon (2026). "The AI Fever Dream: How the Industry Sold a Revolution Nobody Was Ready For." LinkedIn, 26 April 2026.
31. Which? (2025). Super-Complaint: Home and Travel Insurance Claims Handling. London: Which?, September 2025.
32. Zagajewski, Mariusz (2026). "How an Underwriting Workbench Strengthens Regulatory Compliance: The 2026 NAIC AI Bulletin Playbook." Decerto, 21 August 2025; updated 9 May 2026.
33. The underlying research paper, The Aetiology of Boundary Blindness, version 7.5, May 2026, contains the full methodology, all regulatory citations in extended form, the convergence-cluster analysis in detail, and the symptom-to-root-cause mapping appendices. Available on request.