

Verisopic

AI Claims Decision Reconstruction Scrutiny Simulation Briefing

Motor — Low-Value Property Damage
Decision Timestamp: 19 August 2024 · 14:32 UTC

Independent reconstruction of a single AI-assisted claims decision under simulated
ombudsman, regulatory, or litigation review conditions.

Prepared for illustrative purposes
Demonstrating decision-state defensibility testing

February 2026

Simulated example. Not client data.

Verisopic Evidence Standard (VES) v1.0

Document ID: VES-CRS-2026-02-SIM-2024-08

AI Claims Reconstruction Example (Simulated)

Decision Context

This example demonstrates reconstruction of a single AI-assisted claims decision under simulated scrutiny conditions.

Claim Type: Motor — Low-value property damage

Decision: Automated rejection (AI-assisted)

Decision Date: 19 August 2024 · 14:32 UTC

AI Model: Fraud Triage Model v3.2

Human Acceptance: Claims officer approved AI recommendation

Scrutiny Scenario

This reconstruction simulates an ombudsman review occurring approximately 18 months after the original decision.

The purpose is not to assess claim correctness, but to test whether the exact decision state - authority, policy version, system context, and human oversight - can be independently reconstructed to withstand scrutiny.

This example illustrates how evidentiary durability differs from general governance documentation.

The pages that follow demonstrate:

- What governance artefacts existed at the decision timestamp
- Whether authority and policy versions were fixed in time
- Whether reconstruction depends on narrative interpretation
- How integrity verification supports independent review

Under ombudsman, regulatory, or litigation review, inability to independently reconstruct the exact decision state can shift focus from claim merits to governance exposure.

This is a simulated example for demonstration purposes. It does not represent client data.

Reconstruction Findings (Simulated)

Reconstruction Attempt

Using the canonical evidence snapshot aligned to the decision timestamp (19 August 2024 · 14:32 UTC), the following artefacts were identified:

- Governance policy version in force: 1.4
- Fraud threshold applied: >0.78
- Declared responsible authority: Claims Oversight Lead
- AI system registry status: Declared
- Drift classification: No material drift recorded since prior review

Integrity verification confirms that the canonical evidence representation has not been altered since issuance.

Defensibility Assessment

Reconstruction Strength: Conditional

Core governance artefacts were present and structurally verifiable. However, reconstruction revealed areas where evidentiary durability may depend on internal interpretation rather than independently fixed records.

Observed Gaps

- Human acceptance record was not cryptographically bound at decision time
- Responsibility declaration was incomplete at the timestamp
- System state export relied on internal operational logs rather than independently fixed artefacts

These gaps do not imply regulatory breach or decision error. They illustrate where exposure may arise under independent scrutiny.

Purpose of This Example

This simulated reconstruction demonstrates the distinction between:

- Having governance documentation
- Being able to independently evidence the exact decision state months or years later

The following pages contain the integrity-bound Evidence Pack structure associated with this simulated decision snapshot under the **Veriscope Evidence Standard (VES)**, illustrating how integrity-bound artefacts support defensibility analysis.



Veriscope Evidence Standard (VES)

A deterministic evidence format for governance artefacts with independent integrity verification.

VES declaration

Standard: VES

Version: 1.0

Issued (UTC): 2025-01-01T00:00:00.000Z

Issuer: Sample Organisation (Redacted)

Issuer ID: [REDACTED]

Integrity

Hash algorithm: SHA-256

Canonical checksum 0ba836e347...6c3154

Verification

Verification method: Canonical SHA-256 recomputation

Verification class: Deterministic (local or independent third-party)

Verification endpoint: <https://verify.veriscope.com>

Verification confirms integrity only — that the canonical Evidence Pack JSON has not been altered since issuance. It does not certify compliance, regulatory sufficiency, risk classification, or outcomes.

VES Manifest

Identifiers

Issued (UTC): 2025-01-01T00:00:00.000Z
Issuer organisation: Sample Organisation (Redacted)
Issuer organisation ID: [REDACTED]
VES version: 1.0

Pack properties

Mode: sample
VES issuance class: VES-S (Sample)
Evidence pack version: 1.1
AI framed: No
Declared systems: 1

Integrity binding

Hash algorithm: SHA-256

Canonical checksum 0ba836e347d381cf131f25359d19318a5f26bdbcb1b8287a8349e9405b16c3154

Verification

Verification method: Canonical SHA-256 recomputation
Verification class: Deterministic (local or independent third-party)
Verification endpoint: <https://verify.veriscopic.com>

Disclaimers

- Integrity only: verification confirms that the canonical evidence representation has not been altered.
- Not legal advice.
- Not a compliance certification.



Governance Evidence Pack — Sample (Redacted)

Public sample. Demonstrates structure and verification approach. Not client evidence.

Pack metadata

Organisation: Sample Organisation (Redacted)
Organisation ID: [REDACTED]
Generated (UTC): 2025-01-01 00:00:00.000 UTC
Evidence Pack Version: 1.1
Declared decision systems: 1

Integrity

This pack is cryptographically bound to a canonical evidence representation. This sample demonstrates verification mechanics. Certain identifiers may be redacted for safety.

Integrity binding applies exclusively to the canonical Evidence Pack JSON. This PDF is a derivative rendering for human inspection.

Checksum 0ba836e347...6c3154

Executive Summary — Governance Evidence Overview

This Evidence Pack is a cryptographically verifiable export of declared governance artefacts recorded within the Veriscopic platform.

What this pack establishes

- Declared governance artefacts are captured in canonical form.
- Document acceptance evidence is version-bound and time-stamped (where present).
- Organisation audit events provide traceable change history (where present).
- Declared system records are structured and integrity-bound.
- Independent integrity verification is possible via SHA-256 checksum comparison.

Scope and limitations

- This document confirms existence, structure, timing, and integrity only.
- It does not certify regulatory compliance.
- It does not provide legal advice.
- It does not determine EU AI Act risk classification.
- It does not assert model safety, performance, or adequacy.

Verification in sample mode demonstrates structural integrity mechanics. Identifiers may be truncated for safety.

Evidence Index

Evidence coverage (informational, non-legal)

Coverage status: Developing. This indicator reflects whether core governance evidence categories are populated. It is not a compliance certification or legal determination.

Included sections

Governance Snapshot — 2 audit events

Legal Document Acceptance — 3 acceptance records

Decision Systems Registry — 1 declared systems

Responsibility & Accountability Summary — 0 declared responsibility records

Verification Appendix — Canonicalization and hashing notes

Legal Documents Accepted

Document	Version	Accepted (UTC)	Content hash
Platform Terms	1.0	2025-01-02 09:15:00.000 UTC	sample-has...-terms
Privacy Notice	1.0	2025-01-02 09:15:00.000 UTC	sample-hash-privacy
AI Governance Disclosure	1.0	2025-01-02 09:15:00.000 UTC	sample-has...rnance