Internet Engineering Task Force
Internet-Draft
Intended status: Informational
Expires: 7 October 2026

T. Adebayo
O. Apalowo
F. Makanjuola
Veridom Ltd

5 April 2026

OMP Domain Profile: AI Liability Insurance Underwriting and Parametric Claims Evidence
draft-veridom-omp-aiins-00

Abstract

This document defines a domain profile of the Operating Model Protocol (OMP) for AI systems deployed in contexts covered by AI liability insurance policies, including AI performance warranties, AI errors and omissions coverage, and coordinated AI liability structures. The profile -- designated InsureMark -- specifies how OMP's deterministic routing invariant, Watchtower enforcement framework, and three-layer cryptographic integrity architecture generate per-decision Proof-Points that function as objective parametric trigger data for AI liability insurance claims, and provide independently verifiable underwriting evidence that reduces claims ambiguity and supports premium differentiation.

The InsureMark profile addresses the primary gap in current AI liability insurance underwriting: policies are currently issued based on model-level performance assessments, but claims arise at the level of individual AI decisions. No current AI liability insurance product requires or receives per-decision cryptographic evidence. This profile specifies the technical architecture by which OMP Proof-Points close this gap.

The OMP core specification is defined in the Operating Model Protocol Internet-Draft (draft-veridom-omp).

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Copyright Notice

Table of Contents

1. Introduction

AI liability insurance has emerged as a significant market in response to the growing deployment of AI systems in consequential commercial and regulated contexts. AI performance warranties, AI errors and omissions policies, and coordinated AI liability structures now offer coverage for financial losses arising from AI errors, model failures, and AI-generated harms.

Current AI liability insurance products share a structural limitation: they are underwritten at the model level and adjudicated at the claim level, with no per-decision evidence infrastructure connecting the two. When a claim arises, the insured and insurer must reconstruct what the AI system did in the specific interaction that generated the alleged harm -- often weeks after the fact, from logs not designed for forensic use.

This gap produces two material consequences: claims uncertainty (the inability to reconstruct the precise decision state is the primary source of disputed claims and extended settlement timelines) and underwriting imprecision (policies cover the AI system as a whole without differentiating between the materially different liability profiles of fully autonomous decisions versus supervised decisions).

The Operating Model Protocol (OMP) [I-D.veridom-omp] generates a cryptographically sealed Proof-Point for every AI decision, containing the routing outcome, policy compliance flag, confidence scores, Named Accountable Officer identity (where human oversight was applied), RFC 3161 [RFC3161] TimeStampToken, and SHA-256 hash chain per [RFC8785]. These Proof-Points are independently verifiable by any party without access to the operator's or OMP implementer's infrastructure.

This document defines the InsureMark profile: the domain-specific instantiation of OMP for insured AI deployments. The profile specifies how OMP Proof-Points function as parametric trigger data for AI liability insurance claims, and how the Audit Trace schema extensions for this profile enable premium differentiation based on per-decision evidence quality.

Related OMP domain profiles include the EU AI Act Article 12 profile [I-D.veridom-omp-euaia] and the Legal AI Supervision profile [I-D.veridom-omp-legal]. The OMP specification is also archived at [ZENODO-OMP].

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119] [RFC8174].

2. Terminology

This document uses the terminology defined in [I-D.veridom-omp]. In addition:

Parametric Trigger A pre-defined, objectively measurable event whose occurrence automatically initiates the claims assessment process. Under this profile, the policy_compliance_flag = INVALID is the primary parametric trigger.

Coverage Tier A differentiated insurance coverage level based on the OMP routing state: AUTONOMOUS, ASSISTED, or ESCALATED interactions carry materially different liability exposure profiles.

InsureMark Proof-Point An OMP Audit Trace record generated and sealed under the InsureMark profile, containing all fields defined in Section 4.5.

Policy Compliance Flag The VALID, INVALID, or PARTIAL determination produced by WT-AIINS-02. An INVALID value is the primary parametric trigger under this profile.

Claims Evidence Package The self-contained artefact defined in Section 8, producible within 30 seconds for any interaction in the coverage period, containing all information required for insurer claims assessment without access to the operator’s infrastructure.

3. AI Liability Insurance Underwriting: The Evidence Gap

3.1. Current Underwriting Standards

Leading AI liability insurance products assess the AI system’s training data quality, testing methodology, governance documentation, usage scenarios, and compliance with AI management standards such as ISO/IEC 42001 [ISO-42001]. These assessments answer: is this AI system designed and governed to operate correctly? They do not answer: did this AI system operate correctly in the specific interaction under claim?

3.2. The Decision-Level Evidence Gap

When a claim arises, insurer and insured must reconstruct what the AI system did at decision time: what input data was presented, what policy evaluation applied, what the AI recommended, whether a human reviewed it, and whether the record has remained intact. Where this reconstruction is impossible, settlement depends on negotiation rather than evidence -- producing extended timelines, disputed claims, and coverage limits that price in this uncertainty.

3.3. Cyber Insurance as Precedent

By 2022, cyber insurers moved from recommending audit logs to requiring them as conditions of coverage, with premium differentiation for verified control effectiveness [DELINEA-2026]. The actuarial basis is direct: organisations with verified audit trails have lower claims uncertainty and reduced disputed claim rates. AI liability insurance is at an earlier stage of the same trajectory.

3.4. ISO/IEC 42001 as Partial Precedent

ISO/IEC 42001 certification has been adopted by at least one leading AI liability MGA as a basis for premium differentiation. ISO/IEC 42001 certifies the AI governance process. It does not certify that any specific decision was made correctly. InsureMark Proof-Points are the execution-time evidence layer that completes what ISO/IEC 42001 process certification started. See Section 7.

4. OMP InsureMark Profile

4.1. Routing States and Coverage Tier Differentiation

The three OMP routing states create distinct Coverage Tiers:

AUTONOMOUS The AI system determined the outcome without human review. Highest-risk Coverage Tier: the insurer’s exposure is to errors in the AI system’s autonomous judgment with no human mitigation in the decision chain.

ASSISTED A Named Accountable Officer reviewed the AI recommendation before the final outcome. Lower-risk Coverage Tier: a named, accountable human was present in the decision chain. The Proof-Point records officer identity and review decision, enabling liability differentiation and potential subrogation against the officer’s professional indemnity policy.

ESCALATED Mandatory human intervention occurred. Lowest-risk

Coverage Tier: the AI system identified a condition requiring human judgment. The Proof-Point records escalation reason and Named Accountable Officer intervention.

Underwriters SHOULD offer differentiated premiums by Coverage Tier, with AUTONOMOUS carrying the highest rate per interaction and ESCALATED the lowest.

4.2. Named Accountable Officer as Liability Differentiator

For ASSISTED interactions, the Proof-Point records whether the Named Accountable Officer reviewed and approved an AI recommendation before finalisation. This is material to claims assessment (professional judgment versus AI system failure question) and subrogation (the insurer may have rights against the Named Accountable Officer's professional indemnity policy).

Under this profile, operators MUST record Named Accountable Officer identity for all ASSISTED and ESCALATED interactions. The identity MUST be stable throughout the coverage period: the same identifier MUST refer to the same individual.

4.3. Confidence Score Configuration

This profile does not mandate specific thresholds; these are negotiated at policy inception. However: (a) the AUTONOMOUS routing threshold MUST be documented in the Underwriting Evidence Record; (b) any change to this threshold MUST generate a WT-AIINS-03 event and be notified to the insurer; (c) C_p = 0.0 MUST force ESCALATED routing, ensuring policy compliance failures generate a mandatory human intervention record.

4.4. Watchtower Definitions

4.4.1. WT-AIINS-01: Performance Threshold Gate

Trigger: Composite Confidence Score falls below the configured AUTONOMOUS threshold.

Action: FORCE_ASSISTED.

Claims relevance: Documents AI system self-identified uncertainty and Named Accountable Officer review decision, enabling distinction between AI-uncertain-and-human-approved versus AI-autonomous-and-erroneous in claims assessment.

4.4.2. WT-AIINS-02: Policy Compliance Evidence Gate

Trigger: Evaluated at every interaction.

Action: Computes policy_compliance_flag (VALID/INVALID/PARTIAL). INVALID sets C_p to 0.0, forcing ESCALATED routing.

Claims relevance: policy_compliance_flag is the primary parametric trigger field. An INVALID flag records that the AI system's behaviour deviated from the operator's declared governance policy -- the insurance-equivalent of a covered loss event.

4.4.3. WT-AIINS-03: Configuration Change Gate

Trigger: Any change to OMP routing configuration (thresholds, Watchtower definitions, profile version) during an active coverage period.

Action: FORCE_ESCALATED for the triggering interaction, generating a sealed configuration change record.

Claims relevance: Enables insurer to verify that the configuration at the time of an alleged error matches the configuration disclosed at underwriting. Configuration version mismatch may affect coverage.

4.4.4. WT-AIINS-04: Coverage Scope Verification Gate

Trigger: Interaction type, domain, or risk category not included in the operator's declared coverage scope.

Action: FORCE_ESCALATED. The out-of-scope interaction MUST NOT proceed AUTONOMOUS.

Claims relevance: Prevents undisclosed scope expansion from generating covered claims without the insurer's knowledge.

4.4.5. WT-AIINS-05: Anomalous Output Rate Gate

Trigger: Rate of INVALID policy_compliance_flag determinations for a given interaction type exceeds the configured anomaly threshold within a rolling window.

Action: FORCE_ESCALATED for subsequent interactions of the same type, pending human review.

Claims relevance: Creates a sealed record of model degradation, data drift, or adversarial input detection events, supporting post-market monitoring rights under the policy.

4.5. Audit Trace Schema Extensions

The following fields are REQUIRED under the InsureMark profile, in addition to core fields in [I-D.veridom-omp] Section 7:

* insurance_policy_id: string, REQUIRED. Identifier assigned by the insurer or MGA at policy inception.

* coverage_tier: string, REQUIRED. One of: "AUTONOMOUS", "ASSISTED", "ESCALATED". MUST match routing_outcome.

* policy_compliance_flag: string, REQUIRED. One of: "VALID", "INVALID", "PARTIAL". INVALID is the primary parametric trigger.

* parametric_trigger_activated: boolean, REQUIRED. True if policy_compliance_flag is INVALID or composite Confidence Score falls below insured_performance_threshold.

* insured_performance_threshold: number, REQUIRED. Decimal 0.0-1.0. Minimum composite Confidence Score specified in policy terms.

* coverage_period_id: string, REQUIRED. Identifier for the active coverage period.

* interaction_type_declared: string, REQUIRED. Interaction type as declared in the Underwriting Evidence Record.

* named_accountable_officer_id: string, REQUIRED for ASSISTED and ESCALATED; NULL for AUTONOMOUS. Stable identifier consistent with the Named Accountable Officer registry disclosed at underwriting.

* configuration_version: string, REQUIRED. Semantic version of OMP configuration at time of interaction.

* profile_version: string, REQUIRED. MUST be "VERIDOM-INSUREMARK-v1.0".

5. Parametric Trigger Architecture

5.1. Trigger Field Mapping

Primary trigger: policy_compliance_flag = "INVALID" - Policy Compliance Failure Event.

Secondary trigger: composite Confidence Score below insured_performance_threshold - Performance Threshold Breach Event.

Coverage tier classifier: coverage_tier differentiates AUTONOMOUS from ASSISTED liability for claims assessment and premium calculation.

Configuration integrity: configuration_version mismatch between interaction and underwriting disclosure may affect coverage.

5.2. Claims Event Generation

When parametric_trigger_activated is true, the InsureMark adapter MUST: (a) extract the sealed Proof-Point; (b) verify chain integrity by recomputing SHA-256(payload_canonical) against interaction_hash; (c) verify the RFC 3161 [RFC3161] TimeStampToken; (d) generate a Claims Event Record containing the interaction_id, parametric trigger details, coverage_tier, routing_outcome, named_accountable_officer_id, insurance_policy_id, and full sealed Proof-Point; and (e) submit to the insurer's claims intake system within the policy notification window.

The Claims Event Record is self-contained: an insurer with access only to the record and the Timestamp Authority's public key can verify integrity without access to the operator's infrastructure.

5.3. Chain Integrity Verification for Claims

The OMP Merkle chain structure enables completeness verification: a gap between Proof-Points N and N+2 indicates at least one interaction was not logged. Insurers discovering a chain gap may treat it as a policy condition breach or require explanation before claims assessment proceeds. Operators MUST maintain an unbroken Proof-Point chain throughout the coverage period. Operational interruptions MUST be documented in a sealed Chain Gap Record.

6. Premium Differentiation Framework

The InsureMark profile enables two premium differentiation mechanisms:

Tier-based differentiation: Policies differentiate premiums by Coverage Tier based on the actual distribution of AUTONOMOUS, ASSISTED, and ESCALATED interactions. The distribution is computed from the sealed Proof-Point stream provided as the Underwriting Evidence Record at renewal. Because the stream is independently verifiable, insurers can audit it without relying on operator self-reporting.

Evidence quality differentiation: Deployments implementing the full InsureMark profile with an unbroken Proof-Point chain demonstrate higher AI governance evidence quality. Insurers SHOULD offer reduced premiums for verified, complete InsureMark chains, consistent with the cyber insurance precedent of premium differentiation for verified control effectiveness.

The actuarial basis for both mechanisms is the same: deployments with complete, independently verifiable Proof-Point records have lower claims uncertainty. The probability of a disputed claim approaches zero when a sealed, independently verifiable Proof-Point exists for every interaction in the coverage period.

7. Interaction with ISO/IEC 42001

ISO/IEC 42001 certifies the AI governance process. InsureMark Proof-Points prove each specific decision. The two mechanisms are layered and complementary:

* ISO/IEC 42001: organisational-level, annual audit, certifies design-time governance.

* InsureMark: per-decision level, every interaction, proves execution-time compliance.

Insurers that offer premium differentiation for ISO/IEC 42001 certification can extend their framework to include InsureMark as a second, execution-time evidence tier. ISO/IEC 42001 answers: is the AI governance system designed correctly? InsureMark answers: did the AI system operate correctly in this specific interaction?

8. Claims Evidence Package

Upon a covered claim event, the operator MUST produce a Claims Evidence Package containing:

* The sealed InsureMark Proof-Point for the interaction under claim.

* Chain integrity proof: SHA-256 Merkle root for the coverage period window and chain path from the Proof-Point to the window root.

* Timestamp Authority verification: RFC 3161 TimeStampToken verification output from the OMP Reference Validator [OMP-OPEN-CORE].

* Named Accountable Officer record: for ASSISTED and ESCALATED interactions, officer identity, review timestamp, and review decision.

* Configuration record: configuration_version at time of interaction and sealed configuration history from policy inception.

* Coverage scope confirmation: verification that interaction_type_declared matches the Underwriting Evidence Record.

The Claims Evidence Package MUST be producible within the timeframe specified in the policy terms. Implementations SHOULD be capable of generating it within 30 seconds for any single interaction. The package is self-contained: an insurer, MGA, loss adjuster, reinsurer, or court with no access to the operator's infrastructure or the OMP implementer's systems can verify its integrity using only the OMP Reference Validator [OMP-OPEN-CORE] and the Timestamp Authority's public key material.

9. Security Considerations

The security considerations of [I-D.veridom-omp] apply in full.

Insurance fraud: Operators MUST NOT circumvent WT-AIINS-03. Operating the AI system outside the disclosed configuration while generating technically valid Proof-Points is a material breach of policy conditions.

Privacy: The Proof-Point stream may contain personal data subject to GDPR or equivalent legislation. Operators MUST ensure disclosure to insurers is consistent with applicable data protection obligations.

Timestamp Authority compromise: Operators SHOULD use QTSPs listed in an EU Member State trusted list under eIDAS or equivalent national trust framework, as these operate under regulatory supervision with key management requirements that reduce retroactive fabrication risk.

Chain gap manipulation: Deliberate creation of chain gaps to obscure non-compliant interactions is a material breach of policy conditions.

10. IANA Considerations

This document has no IANA actions.

11. References

11.1. Normative References

[I-D.veridom-omp]
Adebayo, T., Apalowo, O., and F. Makanjuola, "Operating Model Protocol (OMP): A Deterministic Decision-Enforcement Protocol with Externalized Proof-of-Integrity", Work in Progress, Internet-Draft, draft-veridom-omp-00, March 2026, <https://datatracker.ietf.org/doc/html/draft-veridom-omp-00>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC3161] Adams, C., Cain, P., Pinkas, D., and R. Zuccherato, "Internet X.509 Public Key Infrastructure Time-Stamp Protocol (TSP)", RFC 3161, DOI 10.17487/RFC3161, August 2001, <https://www.rfc-editor.org/info/rfc3161>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8785] Rundgren, A., Jordan, B., and S. Erdtman, "JSON Canonicalization Scheme (JCS)", RFC 8785, DOI 10.17487/RFC8785, June 2020, <https://www.rfc-editor.org/info/rfc8785>.

11.2. Informative References

[DELINEA-2026]
Delinea, "Cyber Insurance Coverage Requirements for 2026", https://delinea.com/blog/cyber-insurance-coverage-requirements-for-2026, 2026.

[I-D.veridom-omp-euaia]
Adebayo, T., Apalowo, O., and F. Makanjuola, "OMP Domain Profile: EU AI Act Article 12 Logging and Traceability Requirements for High-Risk AI System Operators", Work in Progress, Internet-Draft, draft-veridom-omp-euaia-00, April 2026, <https://datatracker.ietf.org/doc/html/draft-veridom-omp-euaia-00>.

[I-D.veridom-omp-legal]
Adebayo, T., Apalowo, O., and F. Makanjuola, "OMP Domain Profile: Legal AI Supervision Under ABA Model Rule 5.3 and California Senate Bill 574", Work in Progress, Internet-Draft, draft-veridom-omp-legal-00, April 2026, <https://datatracker.ietf.org/doc/html/draft-veridom-omp-legal-00>.

[ISO-42001]
International Organization for Standardization, "ISO/IEC 42001:2023 -- Information technology -- Artificial intelligence -- Management system", 2023.

[OMP-OPEN-CORE]
Veridom Ltd, "OMP Open Core: Reference Validator and Schema Library", Apache 2.0, https://github.com/veridomltd/omp-open-core, 2026.

[ZENODO-OMP]
Adebayo, T., Apalowo, O., and F. Makanjuola, "OMP -- Operating Model Protocol: A Deterministic Routing Invariant for Tamper-Evident AI Decision Accountability in Regulated Industries", Zenodo DOI 10.5281/zenodo.19140948, March 2026.

Authors' Addresses

Tolulope Adebayo
Veridom Ltd
London
United Kingdom
Email: tolulope@veridom.io

Oluropo Apalowo
Veridom Ltd
Awka
Nigeria
Email: ropo@veridom.io

Festus Makanjuola
Veridom Ltd
Toronto
Canada
Email: festus@veridom.io