AIPS-1

Agent Insurance Policy Standard

Working Paper v0.1 — Draft for Comment

A common framework for the issuance, identification and verification of insurance policies covering AI agent activity

Published by	Kadikoy Limited, Bermuda (Reg. 202302362)
Date	6 June 2026
Status	Draft — open for public comment
Version	0.1 — initial release
Companion to	AIS-1 Agent Identity Standard (ais-1.org); AES-1 Agent Execution Standard (aes-1.org); AAS-1 Agent Auditability Standard (aas-1.org); ARS-1 Agentic Remittance Standard (ars-1.org); AHS-1 Agent Health Standard (ahs-1.org)
Contact	info@aiagentsservices.net
Repository	github.com/Kadikoy1/aips-1
Website	aips-1.org
License	Creative Commons CC0 — no rights reserved

V0.1 KEY FEATURES

INTRODUCES	First open standard for machine-verifiable insurance covering AI agent activity
DEFINES	The Policy Certificate — a structured on-chain artefact representing a regulated underlying insurance policy
SPECIFIES	Five required properties (P1–P5) — Issuer Authority, Coverage Specificity, Trigger Determinism, Settlement Path, On-Chain Verifiability
TIERED	Three tiers — Standard Commercial, Captive and Specialised, Sovereign-Backed
COMPOSES	With AIS-1 (identity), AES-1 (execution environment), AAS-1 (audit), AHS-1 (sectoral profiles)
NEUTRAL	Jurisdiction-agnostic; recognised insurance supervisors maintained as a governance artefact
PRESERVES	Existing insurance regulation in full; AIPS-1 is a representation standard, not a substitute

ABSTRACT

Identity tells a counterparty *who* an agent is. Execution tells the counterparty *where* the agent operates and whether the environment is sealed. Neither answers the question that determines whether a counterparty can rationally transact with an autonomous agent: *if something goes wrong, who pays?*

The Agent Insurance Policy Standard (AIPS-1) defines a common, jurisdiction-agnostic framework for the issuance, identification and verification of insurance policies covering AI agent activity. It specifies five core properties — Issuer Authority, Coverage Specificity, Trigger Determinism, Settlement Path and On-Chain Verifiability — that any conformant Policy Certificate must satisfy. The standard sits alongside AIS-1, AES-1 and AAS-1 as the fourth element of an open agent infrastructure stack.

AIPS-1 is a representation standard, not a regulatory substitute. The underlying Policy is governed by the law of the Issuer's jurisdiction; the Policy Certificate is a machine-readable representation of that Policy designed for use at agent speeds. Insurance regulation, supervision and conduct rules are unchanged by this standard.

CONTENTS

1. Motivation and the Illegible Insurance Problem

2. Definitions

- 2.1 Parties and Roles
- 2.2 Certificates and Instruments

3. The AIPS-1 Standard

- 3.1 The Subject of the Standard
- 3.2 Relationship to the Underlying Policy

4. Core Properties (P1–P5)

5. The AIPS-1 Policy Certificate

6. Classification (Tiers)

- 6.1 Tier I — Standard Commercial Coverage
- 6.2 Tier II — Captive and Specialised Coverage
- 6.3 Tier III — Sovereign-Backed Coverage

7. Verifiability and Resolution

8. AIS-1, AES-1 and AAS-1 Binding

9. Comparison with Existing Frameworks

10. Legal and Regulatory Context

11. Security Considerations

12. Implementation Roadmap

13. Request for Comment

14. Authors

App. A — Policy Certificate Worked Example

App. B — Verification Flow

App. C — Mapping to Existing Frameworks

1. Motivation and the Illegible Insurance Problem

AI agents are now executing transactions, contracting with counterparties, holding funds, and acting autonomously across financial, healthcare, commercial and administrative workflows. The volume of agent-mediated activity is growing at machine speed and shows every sign of becoming a structural feature of the economy.

Risk allocation is the missing precondition for agent-to-agent commerce at scale. Identity is necessary but not sufficient. Execution-environment integrity is necessary but not sufficient. A counterparty deciding whether to transact with an autonomous agent — whether the counterparty is itself an agent, an MCP server, an exchange, a hospital, or a regulated institution — needs a credible answer to one further question: *if something goes wrong, who pays?*

Insurance is the conventional answer to that question for human and corporate counterparties. The problem is that traditional insurance instruments are not legible to machines. A policy reference today is a PDF, a paragraph in a master service agreement, or at best an opaque number in a vendor portal. None of these is suitable for a system in which an agent must, at the moment of accepting a transaction, verify that its counterparty is covered, that coverage extends to the specific risk in question, that the policy is currently in force, that the claim trigger is objectively evaluable, and that a defined settlement path exists.

We term this the **Illegible Insurance Problem**. The cover may exist; what is absent is the verifiable representation of that cover in a form an agent can act on in real time. The consequences are already concrete:

- Counterparties cannot rationally calibrate reliance on an agent without knowing what cover stands behind it.
- Hospitals deploying clinical-AI agents cannot evidence indemnification to clinicians, payers or regulators without manual document collection.
- Exchanges and clearing systems cannot accept agent counterparties without bespoke off-protocol diligence on every cover instrument.
- Regulators cannot supervise systemic agent activity without a common evidentiary format for insurance backing.
- Reinsurance and capital-markets risk transfer cannot price agent-risk pools without a portable, structured representation of cover.

Existing frameworks do not close this gap. ACORD specifies data formats for insurance interchange between human and corporate counterparties but assumes account-based intermediation between regulated parties. The IAIS Insurance Core Principles set supervisory standards for insurers but say nothing about machine-readable cover representations. W3C Verifiable Credentials provide a generic credential format but no insurance-specific profile. AIS-1, AES-1 and AAS-1 provide identity, execution-environment and audit primitives respectively, but none addresses the question of cover.

AIPS-1 closes the gap by defining a portable, structured, on-chain representation of a regulated insurance policy — the Policy Certificate — anchored to AIS-1 identities, composable with AES-1 and AAS-1, and consistent with the operating principles of regulated insurance worldwide.

A note on scope. AIPS-1 is positioned as a representation standard rather than a regulatory framework. Insurance regulation is jurisdictionally specific, well-developed, and not amenable to substitution by a technical standard. AIPS-1 sits on top of existing insurance regulation; it does not replace it. The underlying Policy remains governed by the law of the Issuer's jurisdiction. The Certificate is a representation of that Policy designed for machine consumption, not a new legal instrument.

2. Definitions

2.1 Parties and Roles

Term	Definition
`Issuer`	The regulated insurance entity that has bound the underlying Policy and issues the Policy Certificate. Identified by an AIS-1 DID.
`Policyholder`	The natural or legal person, or AIS-1-identified agent, in whose name the Policy is written.
`Insured Agent`	An AI agent identified by an AIS-1 DID whose activity is covered by the Policy. There may be one or more Insured Agents per Policy.
`Insurance Supervisor`	The regulatory authority responsible for licensing and supervising the Issuer in the Issuer's jurisdiction.
`Recognised Jurisdiction`	A jurisdiction whose Insurance Supervisor has been admitted to the Recognised Jurisdiction List maintained under § 10.
`Relying Party`	Any third party — agent, human or system — that verifies a Policy Certificate as a precondition to accepting a transaction or interaction with the Insured Agent.

`Attesting Authority`	An entity that issues a verifiable credential supporting one or more AIPS-1 claims — most importantly, the Issuer's regulator-issued authorisation credential under P1.

2.2 Certificates and Instruments

Term	Definition
`Policy`	The underlying insurance contract, governed by the law of the Issuer's jurisdiction, between the Issuer and the Policyholder. AIPS-1 does not modify or govern the Policy itself.
`Policy Certificate`	The structured on-chain artefact specified in § 5 that represents the Policy in machine-readable form. The atomic unit of AIPS-1.
`Coverage Scope`	The structured machine-readable description of perils covered, perils excluded, limits, deductibles, named insureds and period of cover.
`Trigger`	The predicate condition specified in the Policy Certificate that causes a claim to arise. A Trigger references one or more Evidence Sources and resolves objectively.
`Evidence Source`	A defined, addressable source — oracle, attestation, regulator filing, court order, AES-1 enclave log, or equivalent — referenced by a Trigger.
`Settlement Path`	The defined claims-handling endpoint, notification protocol, response timeline, payout instrument and fallback dispute mechanism specified in the Policy Certificate.
`Status`	The current state of the Policy Certificate, drawn from the enumeration { `Active` , `Suspended` , `Claimed` , `Exhausted` , `Lapsed` }.
`Canonicalisation`	Deterministic serialisation prior to hashing or signing. Reuses the AIS-1 §6.2 / AAS-1 §6.2 primitive: JCS (RFC 8785) default.

3. The AIPS-1 Standard

3.1 The Subject of the Standard

The subject of AIPS-1 is the **Policy Certificate**.

The Policy Certificate is a structured, signed, on-chain artefact that any Relying Party can resolve, verify and act on without insurer disclosure of the underlying policy document. The Certificate is the artefact against which P1–P5 verification is performed.

AIPS-1 is a sibling standard to AIS-1 and AES-1, not a profile. Its subject (the insurance policy) is neither an agent nor an execution environment, so a separate standard is required. The pattern is identical to AES-1's relationship to AIS-1: a co-equal standard addressing a distinct subject, designed to interoperate by cross-reference.

3.2 Relationship to the Underlying Policy

The Policy Certificate is not the Policy. The Policy is a contract under the law of the Issuer's jurisdiction; the Certificate is a structured representation of that Policy designed for machine consumption and independent verification. The Certificate's authority derives from the Policy; the Policy's enforceability derives from the law of its jurisdiction.

This separation matters for three reasons.

Preserving existing legal infrastructure

Insurance regulation is well-developed, jurisdictionally specific, and not amenable to substitution by a technical standard. AIPS-1 does not attempt to define a new legal instrument or to displace existing regulatory frameworks. The Certificate sits on top of insurance regulation rather than alongside it.

Permitting issuance without re-papering

An Issuer in a Recognised Jurisdiction may issue Certificates against existing policy products by configuring its issuance system to emit the structured artefact. The underlying policy wording, regulatory filings and reserving requirements are unaffected. The Issuer's regulatory relationship with its supervisor is unchanged.

Verifiability without disclosure

The Certificate exposes only what a Relying Party needs to verify — Issuer authority, coverage scope, trigger conditions, settlement path, current status — and not the full underlying policy document. This is consistent with how regulated insurance disclosure already operates: the existence and basic terms of cover are commonly disclosed; full policy wording is not.

> **LIMITATION**
>
> The Certificate is a representation, not a substitute. Issuance of an AIPS-1 Policy Certificate does not create insurance coverage where none has been written under the underlying Policy, and does not modify the rights or obligations of any party under the Policy.

4. Core Properties (P1–P5)

Every AIPS-1 Policy Certificate **MUST** satisfy the following five core properties. A Certificate that does not satisfy all five MUST NOT be issued under this standard. The properties are normative and apply to every Certificate of every tier.

The five properties roll up to a single question: *is this policy credibly enforceable?* A Certificate that satisfies all five is one whose existence, scope, trigger conditions, settlement path and current status can be verified by any third party without trusting the Issuer's word for any of them.

Property	What it requires	What it prohibits	Verification method
P1 · Issuer Authority	The Issuer holds a current authorisation under the insurance regulation of a Recognised Jurisdiction, evidenced by a verifiable regulator-issued credential.	Self-insurance representations; unregulated coverage instruments; lapsed authorisations; certificates issued by entities whose supervisor is not on the Recognised Jurisdiction List.	Issuer DID resolution; verification of authorisation credential; cross-check against Recognised Jurisdiction List.
P2 · Coverage Specificity	Coverage Scope expressed as structured fields conforming to the AIPS-1 Coverage Schema — perils, exclusions, limits, deductibles, named insureds, period of cover.	Free-text-only scope; ambiguous limits; undefined perils; reliance on Issuer interpretation.	Schema validation; deterministic parse of all required fields.
P3 · Trigger Determinism	Trigger conditions expressed as predicates evaluable against declared Evidence Sources.	Triggers depending solely on the Issuer's "reasonable opinion," "sole determination" or undefined external conditions.	Predicate evaluation with reference to declared Evidence Sources; static review of Trigger expressions.
P4 · Settlement Path	A claims-handling endpoint URI, notification protocol, maximum response timeline, payout instrument and fallback dispute mechanism.	Issuer-discretionary timelines; undefined fallback; undisclosed handling endpoints; payout instruments not reachable by the Insured Agent or Policyholder.	Endpoint reachability; verification that the fallback mechanism is referenceable and binding.
P5 · On-Chain Verifiability	The current Status of the Certificate retrievable from public chain data without reliance on Issuer disclosure. Status transitions reflected on-chain.	Status maintained solely in the Issuer's internal systems; off-chain status binding on the Relying Party.	Block explorer view; contract read of Status field; review of Status transition history.

> **DRAFTING NOTE · P3**
>
> P3 (Trigger Determinism) is a hard requirement in v0.1 — Issuer-discretionary triggers are excluded at every tier. This is a deliberate choice to push the market toward objectively-evaluable triggers and away from the "reasonable opinion" language common in legacy policy wording. Public comment is invited on whether v0.2 should permit discretion-based triggers at Tier I only, while retaining the hard requirement at Tier II and Tier III.

5. The AIPS-1 Policy Certificate

The Policy Certificate is a structured data record issued on-chain by the Issuer. The Certificate contains the following fields.

Field	Type	Description
`certificateId`	string	Unique identifier for the Certificate. Format: `aips1:{chain}:{address}:{seq}`.
`aipsVersion`	string	Version of AIPS-1 against which this Certificate is issued. v0.1 implementations MUST set this to `"0.1"`.
`issuerDid`	AIS-1 DID	AIS-1 DID of the Issuer.
`issuerAuthorisation`	VC ref	Reference to the Issuer's regulator-issued authorisation credential.
`issuerJurisdiction`	ISO 3166-1	Alpha-2 code of the Issuer's authorising jurisdiction. MUST appear on the Recognised Jurisdiction List.
`policyholderDid`	AIS-1 DID	AIS-1 DID of the Policyholder.
`insuredAgents`	array	Array of AIS-1 DIDs of Insured Agents covered by the Policy.
`tier`	integer	Tier classification per § 6: `1`, `2` or `3`.
`coverageScope`	object	Structured Coverage Scope object conforming to the AIPS-1 Coverage Schema.
`triggers`	array	Array of Trigger predicates with referenced Evidence Sources.
`settlementPath`	object	Claims endpoint URI, notification protocol, response timeline, payout instrument, fallback dispute mechanism.

`periodStart`	ISO 8601	Start of period of cover.
`periodEnd`	ISO 8601	End of period of cover.
`aesConditional`	string	Optional. AES-1 Enclave Certificate identifier. If present, coverage is conditional on the Insured Agent operating within the referenced enclave.
`status`	enum	`Active` · `Suspended` · `Claimed` · `Exhausted` · `Lapsed`.
`statusHistory`	array	Append-only array of Status transitions with timestamps.
`metadataUri`	URI	Optional URI to off-chain extended metadata. IPFS preferred.
`signature`	signature	Issuer signature over the canonicalised Certificate. Structurally identical to AIS-1 / AAS-1 signature object.

Lifecycle

The Policy Certificate progresses through a defined lifecycle. Status transitions MUST be reflected on-chain at the time of transition.

State	Meaning
`Active`	Cover is in force. Initial state on issuance.
`Suspended`	Cover is temporarily withdrawn under terms of the underlying Policy. Relying Parties MUST treat coverage as unavailable.
`Claimed`	A Trigger has been met and a claim has been notified. Settlement Path is in motion.
`Exhausted`	Aggregate limit reached. No further claims may be paid under this Certificate.
`Lapsed`	Period of cover has expired or the Policy has otherwise terminated. Certificate is read-only thereafter.

6. Classification (Tiers)

AIPS-1 defines three tiers of Policy Certificate, distinguished by the category of Issuer and the presence of reinsurance backing. The tier is a property of the Certificate, not of the underlying Policy; it reflects the strength of the standing-behind, not the wording of the cover. The tier MUST be declared on the Certificate and is verifiable by Relying Parties at runtime.

Tier	Issuer category	Reinsurance	Typical use cases

I — Standard Commercial	Licensed commercial insurer in a single Recognised Jurisdiction	Discretionary	Routine agent operational risk; professional indemnity-equivalent cover; errors and omissions; cyber.
II — Captive and Specialised	Captive insurer, mutual, protected cell, segregated account company or specialised insurance vehicle in a Recognised Jurisdiction	Typically present, may be on-chain	Bespoke risk profiles; novel agent activities not adequately served by the commercial market; in-house risk retention by large agent operators.
III — Sovereign-Backed	Tier II Issuer with sovereign or quasi-sovereign reinsurance backstop, guarantee fund participation, or equivalent	Required	Systemic agent activities (settlement infrastructure, market plumbing, critical-sector deployments) where reliance on the policy is itself a source of systemic risk.

6.1 Tier I — Standard Commercial Coverage

Tier I Certificates are issued by licensed commercial insurers operating in a single Recognised Jurisdiction. They cover routine agent operational risks — professional indemnity equivalents, errors and omissions, cyber, general liability — and use existing commercial policy wordings adapted to AIPS-1's structural requirements. Reinsurance is discretionary at Tier I; the Issuer's own balance sheet is the primary capital behind the Certificate.

6.2 Tier II — Captive and Specialised Coverage

Tier II Certificates are issued by captive insurers, mutuals, protected cell companies, segregated account companies and other specialised insurance vehicles. They are typically used for bespoke risk profiles, novel agent activities not adequately served by the commercial market, and large agent operators retaining risk in-house. Reinsurance is typically present and may itself be expressed on-chain as a separate AIPS-1 instrument (deferred to v0.3).

6.3 Tier III — Sovereign-Backed Coverage

Tier III Certificates carry a sovereign or quasi-sovereign reinsurance backstop, guarantee fund participation, or equivalent structural backing beyond the immediate Issuer's balance sheet. They are intended for systemic agent activities — settlement infrastructure, market plumbing, critical-

sector deployments — where reliance on the policy is itself a source of systemic risk and a private balance sheet is structurally insufficient.

The tier hierarchy is structural, not qualitative. A well-written Tier I policy may provide better cover for a given risk than a poorly-written Tier III policy. The tier signals the resilience of the standing-behind, not the quality of the underwriting.

> **DRAFTING NOTE · CAPTIVE MARKETS**
>
> Several jurisdictions operate developed captive insurance markets — Bermuda, Cayman, Vermont, Guernsey, Singapore among them. These jurisdictions are particularly suited as operational venues for Tier II and Tier III Certificates, given the specialised regulatory, capital and structural arrangements available. AIPS-1 does not prefer or privilege any such jurisdiction; suitability is for the Issuer and the Policyholder to determine.

7. Verifiability and Resolution

A Relying Party verifies an AIPS-1 Policy Certificate at the time of relying on the underlying coverage — typically at the moment of accepting a transaction or interaction with the Insured Agent. The verification flow is designed to be runnable by an autonomous agent without human intervention.

```
// AIPS-1 Policy Certificate verification

const issuer  = await ais1.resolve(cert.issuerDid);
const auth    = await vc.resolve(cert.issuerAuthorisation);

assert(auth.verifySignature());
assert(recognisedJurisdictions.includes(cert.issuerJurisdiction));
assert(aips1.validateCoverageSchema(cert.coverageScope));
assert(aips1.evaluableTriggers(cert.triggers));
assert(cert.settlementPath.endpointReachable());
assert(cert.status === 'Active');
assert(new Date() < new Date(cert.periodEnd));

// Tier III — verify reinsurance backing
if (cert.tier === 3) {
  assert(aips1.reinsuranceAttested(cert));
}

// AES-1 conditional — verify enclave currency
if (cert.aesConditional) {
  assert(await aes1.enclaveCurrent(cert.aesConditional));
}

// Scope alignment with the contemplated risk
```

```
assert(aips1.scopeCovers(cert.coverageScope, contemplatedRisk));

// All checks passed — coverage is verifiable and active
return accept(cert);
```

Reference verification implementations — including a JavaScript client library, a smart contract verifier and an MCP server-side verification module — are deferred to v0.2.

8. AIS-1, AES-1 and AAS-1 Binding

AIPS-1 is designed to compose with the rest of the open agent infrastructure stack.

Standard	Subject	Cross-reference with AIPS-1
AIS-1	The agent and its sponsor	Issuer, Policyholder, and Insured Agents are all identified by AIS-1 DIDs. An AIS-1 Agent Bond MAY reference one or more AIPS-1 Policy Certificates in its `metadata_uri`. An AIPS-1 Policy Certificate MUST reference one or more AIS-1 Insured Agent DIDs.
AES-1	The execution environment	Coverage Scope MAY be conditioned on the Insured Agent operating within an AES-1 certified enclave. Where so conditioned, the `aesConditional` field references the AES-1 Enclave Certificate, and Trigger evaluation MAY reference AES-1 enclave logs as Evidence Sources.
AAS-1	Records of agent activity	AAS-1 Class A records emitted by an Insured Agent are admissible as Evidence Source entries under P3 Triggers. Claim notifications under P4 SHOULD emit an AAS-1 Class A record.
AHS-1	Agent in healthcare context	AHS-1 Tier II / III deployments SHOULD reference one or more AIPS-1 Policy Certificates in the AHS Credential Set's `aipsRefs` field, evidencing professional liability or product liability cover.

9. Comparison with Existing Frameworks

Framework	Scope	Gap addressed by AIPS-1
ACORD	Data formats for insurance interchange between humans and corporate counterparties	Assumes account-based intermediation between regulated parties; not designed for agent-speed runtime verification

IAIS Insurance Core Principles	Supervisory standards for insurers	Sets supervisory expectations; says nothing about machine-readable cover representations
W3C Verifiable Credentials	Generic credential format	Provides the underlying credential format for AIPS-1's Issuer authorisation; not insurance-specific
ISO 22301	Business continuity management	Provides general principles for response timelines; AIPS-1's Settlement Path under P4 is consistent with ISO 22301 conventions
Smart contract escrow protocols	Programmatic value lock-and-release	Capable of conditional payout but carry no Issuer identity, no regulator credential, no defined scope vocabulary; AIPS-1 supplies the missing layer
AIS-1	Agent identity	Identity only; AIPS-1 adds insurance cover
AES-1	Execution environment	Environment integrity only; AIPS-1 adds the financial backstop
AIPS-1 (this standard)	Insurance policy covering AI agents	First open standard for portable, machine-verifiable agent insurance

10. Legal and Regulatory Context

AIPS-1 is jurisdiction-agnostic in specification. It does not require, prefer or privilege any particular jurisdiction's insurance regulation. Jurisdictional variation is expressed through the Recognised Jurisdiction List, which is a governance artefact, not a normative element of the standard.

The Recognised Jurisdiction List

A Recognised Jurisdiction is one whose Insurance Supervisor:

- is a member of the International Association of Insurance Supervisors (IAIS);
- operates a licensing regime consistent with the IAIS Insurance Core Principles;
- maintains a public register of authorised insurers; and
- has been admitted to the AIPS-1 Recognised Jurisdiction List under the governance procedure published with the standard.

The List is maintained as a governance artefact and is amendable by published procedure. Inclusion on the List is a procedural matter and does not constitute substantive endorsement of any particular insurer or product authorised within that jurisdiction.

Relationship to Insurance Regulation

AIPS-1 does not replace, modify or interpret insurance regulation in any jurisdiction. The underlying Policy is governed by the law of the Issuer's jurisdiction. The Certificate is a representation of that Policy. Issuance of a Certificate does not create insurance coverage where none has been written under the underlying Policy and does not modify the rights or obligations of any party under the Policy.

AML / CFT Considerations

Where the Insured Agent is engaged in regulated financial activity, the Issuer's existing AML / CFT obligations under POCR (Bermuda), the Money Laundering Regulations (UK), the Bank Secrecy Act (US) or equivalent frameworks continue to apply. AIPS-1 does not modify those obligations. The structured fields in the Certificate are designed to be compatible with risk-based AML / CFT screening and may be used as inputs to the Issuer's customer due diligence processes.

Captive Insurance Jurisdictions

Several jurisdictions — including but not limited to Bermuda, Cayman, Vermont, Guernsey and Singapore — operate developed captive insurance markets. These provide structural and regulatory arrangements (segregated account companies, protected cell companies, incorporated cell companies, sponsored captives) particularly suited as operational venues for Tier II and Tier III Certificates. AIPS-1 is fully compatible with commercial-market issuance under Tier I and does not require captive issuance.

> **CRITICAL LIMITATION**
>
> AIPS-1 does not provide legal advice and does not determine the regulatory status of any particular Policy or Certificate. It provides a standard technical framework and common vocabulary that enables consistent, evidence-based regulatory assessment across jurisdictions.

11. Security Considerations

11.1 Certificate Integrity. Every AIPS-1 Policy Certificate MUST be signed by the Issuer under the Issuer's AIS-1 verification method. Certificates SHOULD include a `prevHash` field referencing the canonical hash of the immediately preceding certificate in the same issuance series, creating a per-Issuer hash chain.

11.2 Issuer Compromise. If an Issuer's signing key is compromised, all Certificates signed by that key after the compromise are at risk. Issuers MUST publish a key-revocation notice through their AIS-1 verification method and SHOULD use HSM-protected keys.

11.3 Status Manipulation. Status changes are on-chain and signed by the Issuer (or an authorised oracle). Off-chain status representations are non-binding. Relying Parties MUST treat on-chain

status as authoritative.

11.4 Oracle Manipulation. Triggers that depend on oracles inherit the trust assumptions of those oracles. Implementations using oracle-dependent triggers SHOULD require multiple independent sources, witness signatures, and a declared dispute-resolution path. Single-oracle triggers are acceptable only where the oracle is itself an AIS-1-identified party operating under a published policy.

11.5 Settlement Path Denial of Service. A claims-handling endpoint that is unreachable defeats P4. Issuers MUST ensure endpoint availability under the declared response timeline. Persistent unreachability SHOULD invoke the declared fallback dispute mechanism.

11.6 Privacy. The Certificate exposes only the fields necessary for verification. Underlying policy documentation, claims correspondence, and beneficial-ownership data are not part of the Certificate. Where additional disclosure is required (e.g. to a regulator), that disclosure occurs through existing regulatory channels.

12. Implementation Roadmap

Phase	Deliverable	Target
0.1 — This document	Specification, P1–P5 properties, Policy Certificate schema, three-tier classification, worked example, public website (aips-1.org)	June 2026
0.2 — Tooling	did:aips1 DID method specification, reference verifier (JavaScript), smart contract verifier, MCP server module, Coverage Schema v1, first Tier I and Tier II Certificate issuances	Q3 2026
0.3 — Reinsurance	On-chain reinsurance attestation framework, Tier III sovereign-backed framework, syndicated coverage structures, first Tier III Certificate	Q4 2026
0.4 — Captive integration	Captive market integration profiles (Bermuda, Cayman, Vermont, Guernsey, Singapore); regulator dashboard specification	Q1 2027
0.5 — Conformance	Conformance suite, test vectors, reference verifier	Q2 2027
1.0 — Standardisation track	Submission to ISO TC 68 / SC 8 liaison and IAIS observer engagement; stable schemas	Q3 2027

13. Request for Comment

AIPS-1 v0.1 is published as a draft for public comment. Feedback is invited from insurance supervisors (IAIS members, national regulators); commercial insurers and reinsurers; captive insurance markets and managers (Bermuda BMA, Cayman CIMA, Vermont DFR, Guernsey GFSC, Singapore MAS); insurance brokers and risk advisers; AI agent developers; legal, regulatory and compliance professionals; standards organisations (ISO TC 68, IAIS, IFRS Foundation); and the wider AI-agent developer community.

Feedback may be submitted via:

- Feedback form: `aips-1.org/#feedback`
- Email: `info@aiagentsservices.net`
- GitHub: `github.com/Kadikoy1/aips-1/issues`

The comment period for v0.1 closes 30 November 2026. A revised draft will be published as v0.2.

Open Questions for Public Comment

- Should P3 permit Issuer-discretionary triggers at Tier I while retaining the hard determinism requirement at Tier II and Tier III?
- Should the Recognised Jurisdiction List be governed by AIPS-1's own governance body, or delegated to an existing standards organisation (e.g. IAIS)?
- Should Trigger predicates be expressed in a defined predicate language (e.g. JSONLogic, Rego) or remain Issuer-defined within a structural envelope?
- What is the appropriate on-chain home for Policy Certificates: a dedicated registry, the Issuer's own contract space, or a multi-chain identifier scheme?
- Should reinsurance attestations be on-chain artefacts in their own right (suggesting an AIPS-2 reinsurance standard) or remain fields within the Policy Certificate?
- How should Certificate Status changes notify subscribed Relying Parties without exposing the Insured Agent's transaction graph?

14. Authors

Author	Kadikoy Limited, Bermuda
Affiliation	BDA Law; BDA AI Agent Services
Companions	AIS-1 Agent Identity Standard (ais-1.org); AES-1 Agent Execution Standard (aes-1.org); AAS-1 Agent Auditability Standard (aas-1.org); ARS-1 Agentic Remittance Standard (ars-1.org); AHS-1 Agent Health Standard (ahs-1.org)
Contact	info@aiagentsservices.net

Website	aips-1.org
Repository	github.com/Kadikoy1/aips-1
License	Creative Commons CC0. No rights reserved. Open for free implementation.

Appendix A — Policy Certificate Worked Example

A Bermuda-licensed captive insurer issues a Tier II Policy Certificate covering an AI clinical decision support agent operating in a US deploying institution. Cover scope is professional indemnity for diagnostic errors. Trigger is an oracle attestation of an upheld malpractice claim. Settlement is via on-chain stablecoin payout within 30 days of trigger satisfaction.


```
{
  "certificateId": "aips1:base:0x4f3edf...:001",
  "aipsVersion": "0.1",
  "issuerDid": "did:ais1:sponsor:bermuda-captive-issuer-example",
  "issuerAuthorisation": "https://creds.bma.bm/insurer/0xabc...",
  "issuerJurisdiction": "BM",
  "policyholderDid": "did:ais1:sponsor:us-hospital-example",
  "insuredAgents": [
    "did:ais1:base:clinical-cds-agent-001"
  ],
  "tier": 2,
  "coverageScope": {
    "perils": ["diagnostic_error", "missed_diagnosis", "incorrect_recommendation"],
    "exclusions": ["intentional_misuse", "off_label_deployment"],
    "perOccurrenceLimit": { "currency": "USD", "value": "1000000" },
    "aggregateLimit":    { "currency": "USD", "value": "5000000" },
    "deductible":        { "currency": "USD", "value": "25000" },
    "namedInsured": "did:ais1:sponsor:us-hospital-example"
  },
  "triggers": [
    {
      "triggerId": "01HZ9D7Y2K3M4N5P6Q7R8S9T0V",
      "predicate": {
        "type": "oracle",
        "sourceRef": "https://oracles.example.org/malpractice-feed",
        "field": "claimUpheldAgainstAgent",
        "operator": "eq",
        "value": true
      },
      "evidenceRequired": ["attestation", "hash_anchor"]
    }
  ],
  "settlementPath": {
    "claimsEndpoint": "https://claims.example-captive.bm/aips1",
    "notificationProtocol": "AIPS-1 Class V notification",
    "responseTimeline": "PT720H",
```

```
    "payoutInstrument": "USDC on Base",
    "fallbackDisputeMechanism": "Bermuda International Court of Arbitration"
  },
  "periodStart": "2026-06-06",
  "periodEnd":   "2027-06-06",
  "aesConditional": "aes1:base:enclave-cleveland-clinic-001",
  "status": "Active",
  "statusHistory": [
    { "state": "Active", "at": "2026-06-06T10:00:00Z" }
  ],
  "metadataUri": "ipfs://Qm...",
  "signature": {
    "alg": "EdDSA",
    "hashAlg": "SHA-256",
    "canonicalisation": "JCS",
    "keyRef": "did:ais1:sponsor:bermuda-captive-issuer-example#key-1",
    "value": "z58dYMy3..."
  }
}
```

Appendix B — Verification Flow

How a Relying Party — counterparty agent, exchange, hospital information governance committee, regulator — verifies an AIPS-1 Policy Certificate before relying on the underlying cover:

1. Receive the Certificate identifier or fetch from the Issuer's on-chain registry.
2. Resolve the Issuer's AIS-1 DID via the AIS-1 §7.1 resolution algorithm.
3. Resolve the Issuer's authorisation credential. Confirm currency and validity. Confirm the issuing Insurance Supervisor appears on the Recognised Jurisdiction List.
4. Validate the Certificate's signature against the Issuer's verification method.
5. Validate that the Certificate satisfies P1–P5 by reference to declared fields.
6. Read on-chain Status. Confirm `Active` and within period of cover.
7. Confirm Coverage Scope aligns with the contemplated risk and that the contemplated Insured Agent is named in `insuredAgents`.
8. Where `aesConditional` is set, confirm via AES-1 that the Insured Agent is currently operating within the referenced enclave.
9. For Tier III, verify the reinsurance attestation.
10. Proceed or decline. Record the verification outcome as an AAS-1 Class A record of type `policy_check`.

Appendix C — Mapping to Existing Frameworks

Framework / requirement	AIPS-1 fields supporting evaluation
IAIS ICP 1–3 (objectives and powers of supervisor)	issuerJurisdiction + Recognised Jurisdiction List
IAIS ICP 4 (licensing)	issuerAuthorisation VC + authorisation credential currency
IAIS ICP 14 (valuation)	coverageScope (limits, deductibles) + reserving disclosed off-chain
IAIS ICP 19 (conduct of business)	settlementPath response timeline + fallback dispute mechanism
IAIS ICP 22 (AML / CFT)	insuredAgents + Policyholder DID (resolves to AIS-1 KYC status)
EU Solvency II disclosure	off-chain via metadataUri; AIPS-1 supplies stable identifiers
ACORD policy interchange	coverageScope schema MAY be ACORD-compatible at v0.2
ISO 22301 incident response timelines	settlementPath.responseTimeline