AIVS-1

Agent Inference Value Standard

Working Paper v0.1 — Draft for Comment

An open standard for recording, proving and valuing the cognitive work autonomous agents deliver — in which occurrence is attested and value is revealed by parties other than the producer.

Published by	Kadikoy Limited, Bermuda (Reg. 202302362)
Status	Draft — open for public comment
Version	0.1 — initial release
Companions	AIS-1 Agent Identity Standard (ais-1.org); AAS-1 Agent Auditability Standard (aas-1.org); AES-1 Agent Execution Standard (aes-1.org); AIPS-1 Agent Insurance Policy Standard (aips-1.org); ARS-1 Agentic Remittance Standard (ars-1.org); ADFACS-1 Agent Dispute Facilitation Standard (adfacs-1.org)
Contact	info@aiagentsservices.net
Repository	github.com/Kadikoy1/aivs-1
Website	aivs-1.org
Licence	Creative Commons CC0 — no rights reserved

V0.1 KEY FEATURES

INTRODUCES	First open standard for the unit of agent inference value
DEFINES	The Inference Value Record (IVR) — a signed, identity-bound record of one inference, its occurrence proof, and the value committed against it
SEPARATES	Occurrence from value. What happened is attested; what it is worth is revealed
GRADES	Occurrence across four attestation tiers (A0–A3) and transferability across three (T0–T2)
QUALIFIES	Attributions by whether the commitment is escapable — irrevocability and an independent trigger, not the identity of the committer
COMPOSES	With AIS-1 (identity), AAS-1 (audit), AES-1 (execution), AIPS-1 (insurance), ARMS-1 (receivables)
DISCLAIMS	Any claim to measure benefit. § 5.4 states plainly what revealed value is not

Abstract

Agent work is priced today in the currency of its inputs: tokens processed, GPU-seconds consumed, FLOPs burned. Those are the substrate, not the output. A determination that prevents a misdiagnosis and a determination that summarises a memo may cost the same to produce and be worth nothing like the same amount. Cost is a property of the producer, and the producer's cost reveals nothing about worth. We term this the **Cost-Proxy Problem**.

The Agent Inference Value Standard (AIVS-1) defines the unit in which the cognitive work an autonomous agent delivers is recorded, proved and valued. That unit is the **Inference Value Record**. The standard rests on a single separation: what happened is *attested*, graded by the strength of the available proof; what it is worth is *revealed*, taken from what a party other than the producer has irrevocably committed contingent on the inference being right.

AIVS-1 makes no claim to measure the benefit an inference confers. What third parties commit against is exposure to the consequences of error, which is closer to materiality than to worth. That limit, and four others, are stated in the body of the standard rather than buried, because the standard's credibility rests on the limits being as legible as the mechanism.

Published as a working paper for public comment, CC0, with a JSON Schema, validated worked examples and a reference verification flow.

Contents

IN ESSENCE

One.	The unit is the **Inference Value Record**. One inference, one record: who produced it, what it was, graded proof that it happened, what independent parties committed against it, and a signature. Minted inside a certified execution enclave so that issuance itself is trustless.
Two.	**Occurrence and value are different questions.** No ledger can verify a frontier-model inference by re-execution, so occurrence is established by attestation and graded A0 to A3. Value is not attested at all — it is read off what somebody else put at risk.
Three.	**The producer never values its own inference.** An attribution counts only if the committer is not the producer, the commitment is irrevocable, and forfeiture is determined by an independent party. A sponsor standing behind its own agent qualifies — that is the ordinary case, not an abuse.
Four.	**Nothing is collapsed into a scalar.** A premium is not a value; it carries the insurer's cost of capital and the market cycle. The record holds the components and leaves the arithmetic to the reader.
Five.	**It is candid about reach.** Revealed value is a floor rather than a measure, exposure rather than upside, and exists only where somebody happened to commit. Most inferences attract nothing and record nothing. § 5.4 says so.

1. Motivation and the Cost-Proxy Problem

Autonomous agents now form judgments that carry consequences: screening a counterparty, pricing a risk, allocating capital, reading a scan, approving a payment. The volume is growing at machine speed and shows every sign of becoming a structural feature of the economy.

None of it has a unit.

Where agent work is priced at all, it is priced by what it consumed — tokens, GPU-seconds, model calls. That is the substrate on which the work ran, not the work. Two inferences of identical cost may differ in consequence by six orders of magnitude, and the producer's cost schedule contains no information about which is which. We term this the **Cost-Proxy Problem**: the thing being measured is an input, and it is measured because it is easy to measure rather than because it is informative.

The consequences are already concrete. An inference cannot be booked, because there is no unit in which to book it. It cannot be audited as a discrete item. An underwriter asked to cover agent activity has no denominator against which to price. A counterparty deciding how far to rely on an agent's determination has nothing to read except the producer's own account of its quality. And because no comparable record exists across producers, no portfolio view of agent activity is possible at all.

The obvious alternative — asking the producer what its inference was worth — fails for the reason all self-assessment fails. The producer is the party with both the least information about consequence and the strongest interest in the answer.

AIVS-1 therefore takes its signal from a third position: the party who has put something at risk on the inference being correct. A commitment somebody has irrevocably made, forfeitable on somebody else's determination, is a fact about the inference that survives contact with an adversary. It is not the whole of value, and § 5.4 is explicit about the gap. It is, however, observable, external, and hard to fake — which is more than any of the alternatives can claim.

2. Definitions

2.1 Parties and roles

Term	Definition
AI Agent	A software system that perceives its environment, makes decisions, and takes actions to achieve goals. As defined in AIS-1.
Producer	The agent that performed the inference, identified by an AIS-1 DID. The party whose cost does not reveal value.
Committer	A party other than the Producer that has placed value at risk contingent on the inference's correctness or outcome. Identified by an AIS-1 DID.
Trigger Authority	The party that determines whether forfeiture or loss occurs. Must be neither the Committer nor the Producer.
Relying Party	Any third party reading an IVR as an input to a decision — a counterparty, an underwriter, an auditor, a regulator, a pool administrator.
Minting Enclave	The AES-1 certified contract that validates an attestation and issues the IVR. Its AES-1 tier bounds the record's transferability (§ 8).

2.2 Records and commitments

Term	Definition
Inference	The act of a model producing a result — a determination, completion, classification, or judgment — from an input. The delivered cognitive work. Distinct from compute (§ 4).
Compute	The input substrate: FLOPs, GPU-time, tokens processed. The resource consumed to produce inference. Not the unit AIVS-1 measures.
Inference Value	The magnitude of consequence parties other than the Producer commit contingent on an inference's correctness or outcome. **Used throughout as a term of art.** It is not a measure of the benefit an inference confers when correct; § 5.4 states the difference.
Inference Value Record (IVR)	The atomic unit of AIVS-1. A signed, AIS-1-bound document attesting an inference, its occurrence tier, and the value committed against it. An accounting record always; transferable only by tier.
Attribution	A single recorded instance of a Committer placing value at risk against an inference, citing a signal class, its components, the Committer, and evidence.
Qualifying Attribution	An Attribution satisfying § 5.3: the Committer is not the Producer, the commitment is irrevocable, and the Trigger Authority is independent of both. Only Qualifying Attributions count toward transfer tiering.
Irrevocability	The Committer cannot unilaterally recover the committed amount. Segregated, escrowed, or locked beyond the Committer's reach.
Commitment	A single act of placing value at risk, identified by `commitmentId`, which may be apportioned across more than one inference (§ 5.5).
Attestation Tier	The graded strength of proof that an inference occurred. A0 self-attested, A1 TEE-attested, A2 optimistic, A3 zero-knowledge (§ 6).
Transfer Tier	The derived, read-only property determining whether a record may be assigned or pooled. T0 to T2 (§ 8).
Subject Inference	The inference whose value an IVR records. Typically referenced as an AAS-1 record — a Class A action, or for independent findings over a population, a Class D determination.
Canonicalisation	Deterministic serialisation prior to hashing or signing. AIVS-1 specifies JCS (RFC 8785), matching the stack.

> **DRAFTING NOTE · "BOND"**
>
> An **AIS-1 bond** is the identity link between an agent and its sponsor and contains no capital. The AIVS-1 signal class for capital at risk is `posted_stake`. Earlier drafts used `bonded_stake`, which collided with AIS-1 terminology inside a single stack; implementations should treat it as a deprecated alias. "Bond" is reserved throughout for the AIS-1 identity artefact.

3. The AIVS-1 Standard

3.1 The Inference Value Record

AIVS-1 defines the Inference Value Record as the fundamental unit of agent inference value. An IVR is a structured, signed, identity-bound document capturing a single inference, the graded proof that it occurred, and the value independently committed against it.

Every IVR contains five primitives: a **producer identity** (via AIS-1); a **subject-inference reference** (typically an AAS-1 record); an **occurrence attestation** at a stated tier; one or more **revealed-value attributions**; and a **signature**. The record is minted inside an AES-1 certified enclave so that issuance itself is trustless — no operator can mint an IVR that the architecture does not permit.

3.2 What AIVS-1 does not do

The standard is deliberately narrow, and three boundaries matter more than the rest.

It does not create obligations. An IVR records that a party committed capital against an inference. It does not oblige anyone to pay anything. There is no payment stream in an IVR.

It does not price. AIVS-1 records what others revealed. It computes nothing, values nothing, and takes no view on whether any attribution is economically sound.

It does not measure benefit. The unit captures exposure to the consequences of error. Where that differs from worth — and it frequently does — the standard captures the former. See § 5.4.

4. Inference and Compute

The distinction is load-bearing and worth stating precisely. Compute is consumed; inference is delivered. Compute is metered by the infrastructure provider and is a cost to the producer; inference is the result on which other parties act.

Compute may be disclosed on a record as substrate context. It must never function as a price. The distinction is not merely definitional: a standard that permitted cost disclosure to serve as a value figure would reintroduce the Cost-Proxy Problem through a side door, and every incentive in the market would push implementers through it.

5. Revealed Value

5.1 The anchor

> **REVEALED VALUE**
>
> The magnitude of consequence a party *other than the producer* commits — capital at risk, liability assumed, or reliance placed — contingent on the inference's correctness or outcome.

Producers know their costs. Buyers know their own situation. Neither knows what an inference is worth and both have reason to say otherwise. AIVS-1 takes its signal from the party who has put something at risk on the inference being correct.

Four signal classes are recognised in v0.1: **written cover** (an insurance policy), **posted stake** (capital forfeitable if the inference proves wrong), **downstream reliance** (an action of stated magnitude taken on the strength of the inference), and **collateral** (assets locked against the outcome). Appendix C sets out what evidences each.

5.2 Written cover

Written cover leads because it is the cleanest instance and the one the standard can most readily demonstrate. It is also the one most easily misread, and the misreading needs stating before the mechanism.

> **A PREMIUM IS NOT A VALUE**
>
> Approximately, `premium ≈ P(loss) × E[loss | loss] × loading`.
>
> The expected-loss term measures harm if the inference is *wrong*, not benefit if it is *right*. The loading term reflects the insurer's cost of capital, expense ratio, correlation with its existing book, and position in the market cycle — none of which is a property of the inference. Identical inferences written in a hard and a soft market attract materially different premiums for reasons wholly external to them.

AIVS-1 therefore records the policy's **components** and collapses none of them: premium, per-occurrence limit, aggregate limit, deductible and period, read from the AIPS-1 Policy Certificate. The limit is the more stable exposure figure. The premium is the market's price for carrying that exposure at a moment in time. Rate on line — premium divided by limit — is derivable by the reader, which keeps the cycle visible rather than buried inside a single number.

Insurability composes directly with the rest of the stack. The insurable object is typically an AAS-1 Class D determination. The minting enclave's AES-1 tier is itself an underwriting input: a Tier III Sovereign Enclave determination is cheaper to cover than a self-attested one, so certification and revealed value reinforce each other.

AIVS-1 does *not* depend on insurance. Where no underwriter is present, posted stake, downstream reliance and collateral are same-class signals. Insurability is the example, not a load-bearing dependency of the definition.

5.3 Qualification and related parties

An attribution counts toward transfer tiering only if it **qualifies**. Three conditions, all required.

Condition	Requirement	Field
Independence	The Committer is not the Producer. A producer cannot value its own inference. Absolute; no exceptions.	`committerRef`
Irrevocability	The Committer cannot unilaterally recover the amount. Segregated, escrowed, locked on-chain, or otherwise beyond reach.	`irrevocable`
Independent trigger	Forfeiture or loss is determined by a party other than the Committer and the Producer — an AAS-1 Class D determination, an ADFACS-1 determination, an oracle, a counterparty claim, or a court.	`triggerAuthority`

Capital held in a segregated account and forfeitable on somebody else's determination is real exposure whoever posted it. A book entry the Committer can reverse next Tuesday is not exposure at all — it is a statement of intent wearing the costume of a commitment. This is the distinction the law already draws between a genuine trust or security interest and a mere accrual, and it does the work that a rule about the *identity* of the committer cannot.

Non-qualifying attributions are recorded and disclosed. They count toward nothing.

Related parties

`relatedParty` is true where Committer and Producer resolve to the same AIS-1 sponsor. It is a **disclosure carrying a tier ceiling**, not an exclusion.

Attribution	May reach T1	May reach T2
Qualifying, related party	Yes	No — an unrelated qualifying attribution is additionally required
Qualifying, unrelated	Yes	Yes
Non-qualifying	No	No

A sponsor putting capital behind its own agent is the ordinary case, not an abuse of the standard. It is what the AIS-1 accountability chain contemplates: the sponsor is the responsible person, permanently bonded to the agent and answerable for it. A rule counting only commitments from strangers would return nothing in the common case and would penalise precisely the sponsors behaving most responsibly.

The risk the ceiling addresses is **circular attribution**. Where an underwriting agent's own inferences are valued by the premiums it writes, and those premiums are the signal for other agents' inferences, attributions can circulate within a sponsor group without external capital entering the system. That risk bites at admission to a pool, which is what the T2 ceiling protects. It does not bite on a sponsor standing behind its own agent.

5.4 What revealed value is not

The credibility of this standard rests on this subsection as much as on the mechanism it qualifies. Each limit below is real, known, and not cured by anything elsewhere in the specification.

FIVE LIMITS

1 · A floor, not a measure. An attribution records what one party was prepared to commit. It is a lower bound on consequence, not a valuation. Nothing here asserts that an inference is *worth* the sum attributed to it.

2 · Exposure, not upside. Third parties commit against the consequences of an inference being wrong. A correct inference with large upside and an incorrect one with large downside attract signals of similar magnitude. What the record captures is closer to **materiality** than to worth — which is exactly what an auditor, an underwriter or a supervisor needs, and is not a price.

3 · Censored. A signal exists only where somebody happened to commit. The great majority of agent inferences attract nothing and resolve to T0. Any statistic computed over IVRs is computed over the insured, staked and relied-upon minority — a selected population, not a representative one.

4 · Silent on outcome. An IVR records what was committed *ex ante*. v0.1 records nothing about whether the inference held. Until that loop closes the standard produces unclosed valuations and cannot support a track record.

5 · Evidence, not re-execution. In every attestation tier the heavy inference happens off-chain and only the *evidence* is verified on-chain. AIVS-1 does not claim to make proof of inference cheap or solved.

These limits are stated in the body rather than an appendix deliberately. "Inference value" is defined in § 2.2 as a term of art precisely because its ordinary meaning would overclaim, and a stipulated definition is only honest where the gap it stipulates around is legible.

5.5 Apportionment across inferences

Where one commitment is contingent on more than one inference, it must not be attributed in full to each. Agent A screens a counterparty; agent B, differently sponsored, releases funds in reliance on A's conclusion. The commitment is contingent on A's screening *and* on B's decision to rely on it. Attributing it wholly to either misstates that agent's record; attributing it wholly to both double-counts the same money across the system.

The Committer apportions. The party with capital at risk knows why it committed and which inferences it relied on, and states the split. Each attribution carries a `commitmentId` identifying the underlying commitment and a `share` stating the proportion attributed to this inference. Absent `share`, the whole commitment is attributed to that inference.

One rule is enforced, and it is arithmetic:

APPORTIONMENT INVARIANT

The sum of `share` across all attributions citing the same `commitmentId` MUST NOT exceed 1.

This is checkable in tooling and sufficient to stop the same commitment appearing as value on two agents' records. It is *not* sufficient to make the split correct. A Committer may apportion self-servingly or arbitrarily and AIVS-1 has no way to detect it.

The alternative — allocating a commitment across contributors by each one's marginal contribution — is well developed in cooperative game theory and would give a principled answer. It is not available to this standard. A marginal-contribution allocation is *computed* by an operator against a counterfactual, whereas every signal class here requires that an external party actually committed something. Adopting it would abandon the premise. Committer apportionment is therefore the honest boundary of the revealed-value approach: internally consistent, externally unverified, and where a principled allocation is genuinely required, the problem lies outside this standard.

6. The Attestation Ladder

Beneath revealed value sits the plainer question of whether the inference happened at all. No ledger can verify a frontier-model inference by re-execution — re-running costs as much as producing, and inference is non-deterministic across hardware, batch size and numerical precision. Occurrence is therefore established by attestation, graded:

Tier	Mechanism	Trust placed in	Status
A0 — Self	The producer signs a claim that the inference occurred.	The producer. Weak.	Specified
A1 — TEE	A trusted execution environment (SGX, Nitro, confidential-compute GPU) emits a hardware-signed attestation that the model ran in a sealed enclave.	Hardware and the manufacturer's attestation service.	Specified; practical now
A2 — Optimistic	The claim is assumed honest and is challengeable; a challenger forces re-execution under a fraud proof, with the producer's stake slashed if the claim was false.	Economics and the credible threat of challenge.	Reserved; v0.2
A3 — Zero-knowledge	A succinct zero-knowledge proof that a specific model ran on a specific input, verified on-chain for negligible cost.	Mathematics only.	Reserved; roadmap

A1 is the realistic v0.1 occurrence tier and the mechanism the wider stack already contemplates: AAS-1 recognises a TEE-quote evidence type, and AES-1 enclaves can require enclave-attested execution.

The occurrence object carries a `modelIdentityHash` binding the attestation to a specific model identity. An attestation that *some* model ran is not an attestation that *this* model ran, and without the binding the tier claims more than the evidence supports.

The standard is attestation-source-agnostic: the value layer is unchanged regardless of which rung proved the inference. The ladder grades the *strength* of the occurrence proof, which in turn gates transferability (§ 8).

> **A NOTE ON HONESTY**
>
> In every tier the heavy inference happens off-chain and only the evidence is verified on-chain. That is the unavoidable shape of the problem. AIVS-1 does not claim to make proof of inference cheap or solved; it grades how strong the available proof is and prices that into the record's transferability. Neither A2 nor A3 is normatively specified in v0.1 — A2 depends on challenge economics not yet demonstrated in production for inference workloads, and a fraud proof nobody can afford to bring is not a fraud proof.

7. The Inference Value Record — Fields

7.1 Core fields

Attribute	Description	Req.
aivs	Standard version. `"0.1"`.	Yes
recordId	ULID or UUID, unique within the producer.	Yes
producerRef	AIS-1 DID of the agent that produced the inference.	Yes
subjectRef	URI of the subject inference — typically an AAS-1 Class A or Class D record.	Yes
occurrence	Occurrence object: `tier` (A0–A3), `evidenceRef`, `modelIdentityHash`. See § 6.	Yes
mintingEnclaveRef	AES-1 Enclave Certificate reference of the contract minting this IVR.	Yes
timestamp	RFC 3339 timestamp of issuance, with explicit offset. See § 7.3.	Yes
aas1RecordRef	URI of the AAS-1 Class A record emitted at IVR issuance.	Yes
signature	Signature object over the canonicalised record. Structurally identical to AAS-1 § 6.3.	Yes

7.2 Value fields

Attribute	Description	Req.
revealedValue	Array of attributions. At least one required.	Yes
…[].signal	`insurance` | `posted_stake` | `downstream_reliance` | `collateral`.	Yes
…[].components	The commitment as its constituent figures, never as a single derived scalar. For `insurance`: premium, per-occurrence limit, aggregate limit, deductible, period. See § 5.2 and App. C.	Yes
…[].committerRef	AIS-1 DID of the committing party. MUST NOT equal `producerRef`.	Yes
…[].irrevocable	Whether the committer can unilaterally recover the amount. Required condition of qualification (§ 5.3).	Yes
…[].triggerAuthority	AIS-1 DID of the party determining forfeiture. MUST NOT be the committer or the producer. Required where `relatedParty` is true; recommended always.	Cond.
…[].relatedParty	Whether committer and producer resolve to the same AIS-1 sponsor. Disclosed on the record; carries a T2 ceiling (§ 5.3).	Yes
…[].commitmentId	Stable identifier for the underlying commitment. Required where `share` is present.	Cond.
…[].share	Proportion of the commitment attributed to this inference. Absent means 1. Subject to the § 5.5 invariant.	Optional
…[].evidenceRef	URI to on-chain or attested evidence of the commitment.	Yes
valueClass	Capability or quality class of the inference, for comparability. Self-declared and AIS-1-bonded, or independently benchmarked. See § 13, question 1.	Optional
transferTier	Computed transfer tier (T0–T2) per § 8. Read-only; verifiers MUST recompute and MUST NOT trust this field.	Derived

7.3 Canonicalisation and signature

Records are canonicalised with JCS (RFC 8785) and hashed with SHA-256 by default. The signature object is structurally identical to AAS-1 § 6.3, so verification tooling is shared across the stack.

Two implementation constraints follow, and both are traps in practice. Monetary amounts are decimal **strings**, not JSON numbers: a float has no canonical serialisation and would produce different bytes, and therefore different hashes, for the same value across implementations. Timestamps are RFC 3339 — a profile of ISO 8601 that pins a single representation of an instant, always with an explicit offset or `Z` — because ISO 8601 alone permits many representations of the same moment and canonicalisation would not be deterministic.

8. Transfer Tiers

An IVR exists as an accounting record in all cases. It can always be booked, audited and reported. Whether it can be **assigned** or **pooled** is a separate property, derived from the strength of the occurrence proof, the AES-1 tier of the minting enclave, and the set of qualifying attributions.

Tier	Requires	Effect
T0 — Record only	A0 occurrence, or no qualifying attribution, or an uncertified minting enclave.	Valid accounting entry. Not assignable.
T1 — Assignable	A1 or above; at least one qualifying attribution; AES-1 Tier II Verified Enclave.	Assignable to AIS-1-identified counterparties. The attribution travels with the record.
T2 — Poolable	A1 or above; at least one qualifying *unrelated* attribution of class `insurance` or `posted_stake`; AES-1 Tier III Sovereign Enclave.	Admissible to a pool of agent receivables as a valuation and eligibility input.

8.1 Derivation

```
function qualifies(a, producerRef) {
  return a.committerRef    !== producerRef
      && a.irrevocable      === true
      && !!a.triggerAuthority
      && a.triggerAuthority !== a.committerRef
      && a.triggerAuthority !== producerRef;
}

function deriveTransferTier(ivr, enclaveTier) {
  if (ivr.occurrence.tier === 'A0') return 'T0';
  if (!enclaveTier)                 return 'T0';

  const q = ivr.revealedValue.filter(a => qualifies(a, ivr.producerRef));
  if (q.length === 0)               return 'T0';

  const capital = a => a.signal === 'insurance' || a.signal === 'posted_stake';
  const unrelatedCapital = q.some(a => !a.relatedParty && capital(a));

  if (enclaveTier === 'III' && unrelatedCapital)        return 'T2';
  if (enclaveTier === 'II'  || enclaveTier === 'III')   return 'T1';
  return 'T0';
}
```

8.2 What a transfer tier does not mean

Three boundaries, each of which a practitioner will otherwise read the wrong way round.

> **AN IVR IS NOT FUNGIBLE, AT ANY TIER**
>
> Each record is a distinct inference by a distinct producer with distinct evidence at a distinct attestation tier. T2 makes a record *admissible to a pool*; it does not make records interchangeable. Where fungibility is required it is created at the wrapper — a participation interest in a pool may be fungible while nothing underneath it is, which is the ordinary structure of asset-backed finance. All homogenising work therefore falls on `valueClass`, and while that field remains self-declared no pool can rely on it.

> **AN IVR IS NOT A RECEIVABLE**
>
> The record states that a party committed capital against an inference. It creates no obligation on anyone to pay anything. There is no payment stream in an IVR and therefore nothing in it to securitise. Where agent work generates a payment obligation, *that obligation is the asset* and ARMS-1 is the standard that carries it. AIVS-1's role in such a structure is valuation and eligibility criteria — closer to an eligibility schedule than to collateral.

> **AN IVR IS NOT A SECURITY, AND A POOL INTEREST MAY WELL BE**
>
> The valuation layer and the instrument layer must not be conflated. An interest in a pool of agent receivables, sold to investors in expectation of a return generated by the producer's activity, is analysed on its own terms in the relevant jurisdiction and will frequently be a security. Nothing in AIVS-1 is intended to affect that analysis in either direction, and nothing in it should be read as a view on the characterisation of any instrument.

9. Composition with the Open Stack

Standard	Layer	Relationship to AIVS-1
AIS-1	Identity	Who produced the inference, who committed value against it, and who determines forfeiture. Shared sponsorship between producer and committer is what `relatedParty` detects.
AAS-1	Auditability	What the agent did. The subject inference is normally a Class A action record, or a Class D determination. Issuing an IVR itself emits a Class A record.
AES-1	Execution	Where the IVR is minted. Issuance inside a certified enclave means no operator can mint a record the architecture does not permit; the enclave tier bounds transferability.
AIPS-1	Insurance	The mechanism behind the lead signal. Where the class is `insurance`, the evidence is a Policy Certificate and AIVS-1 reads its coverage scope as the attribution's components.
ARMS-1	Receivables	Where the instrument lives. AIVS-1 values an inference and creates no payment obligation; ARMS-1 carries the receivable and AIVS-1 supplies the valuation and eligibility input.
ARS-1	Remittance	How consideration settles once owed.
ADFACS-1	Disputes	What happens when an attribution is contested. An IVR and its evidence references are admissible components of a Canonical Dispute Record in Native mode.

10. Legal and Regulatory Context

AIVS-1 is a recording standard. It creates no instrument, no obligation and no entitlement, and it does not affect the characterisation of anything under any jurisdiction's law.

Accounting. Whether inference activity is recognised, and on what measurement basis, is a matter for the applicable financial reporting framework. AIVS-1 supplies a structured record that such a framework could take as an input. It does not assert that an IVR is an asset.

Securities. See § 8.2. The valuation layer and the instrument layer are distinct, and the standard takes no position on the characterisation of any pool interest.

Insurance. Where an attribution cites a policy, the underlying policy is governed by the law of the issuer's jurisdiction and by AIPS-1's representation framework. AIVS-1 reads the certificate; it does not modify cover, create cover, or affect any party's rights under a policy.

Tax. Out of scope. A base built on revealed value would be unworkable — the population is censored (§ 5.4) and the signal is cycle-sensitive (§ 5.2) — and the administrability considerations that govern tax design point in a different direction from the ones that govern this standard.

> **CRITICAL LIMITATION**
>
> AIVS-1 does not provide legal, accounting, actuarial or investment advice, and does not determine the regulatory status of any record, pool or instrument. It provides a technical framework and a common vocabulary enabling consistent, evidence-based assessment.

11. Security Considerations

11.1 Self-valuation. The prohibition on `committerRef === producerRef` is the standard's first line of defence and is absolute. Implementations must enforce it at issuance, not only at verification.

11.2 Escapable commitments. A commitment the committer can revoke, or whose forfeiture the committer itself determines, is not exposure. Both conditions in § 5.3 must be checked; either alone is insufficient.

11.3 Circular attribution. Where sponsors are shared across producers and committers, value can circulate without external capital entering the system. Disclosure via `relatedParty` plus the T2 ceiling is the v0.1 control. Sponsor-graph analysis across records is deferred.

11.4 Over-attribution. Without the § 5.5 invariant, one commitment can appear as value on many records. Registries and verifiers must accumulate shares per `commitmentId` across the whole record set, not per record.

11.5 Tier inflation. `transferTier` is derived. An issuer stating a tier it has not earned is detectable only if verifiers recompute. They must.

11.6 Attestation-tier overclaim. A stated occurrence tier is not a proved one. Verifiers must confirm that the cited evidence supports the tier claimed, including that `modelIdentityHash` binds the attestation to the model asserted.

11.7 Evidence availability. An `evidenceRef` that ceases to resolve degrades the record. Evidence should be content-addressed or anchored where the commitment is expected to outlive its host.

12. Implementation Roadmap

Phase	Deliverable	Target
0.1 — This document	Specification, IVR JSON Schema, validated worked examples, verification flow, website and repository.	Published
0.2 — Outcome and A2	The Outcome Record: an immutable follow-on class recording whether the inference held, whether cover responded, whether reliance was vindicated. Staked-against-realised across a producer's history is what makes the record a track record rather than a snapshot. Plus the A2 optimistic tier and a reference minting-enclave contract.	Q3
0.3 — Comparability	Value-class benchmarking framework. Transfer-tier registry. Eligibility-criteria profile for admission of T2 records to a receivables pool, drafted against ARMS-1. Sponsor-graph circularity detection.	Q4
1.0 — Standardisation	A3 zero-knowledge attestation profile. Conformance suite and reference verifier. Public IVR registry. Engagement with accounting-standard setters on recognition and measurement of agent inference activity.	2027

13. Request for Comment

AIVS-1 v0.1 is published as a draft for public comment. Feedback is invited from AI agent developers and infrastructure builders; insurers, reinsurers and actuaries; accounting and audit professionals; economists working on agent pricing and coalition value; structured-finance and securities practitioners; legal, regulatory and compliance professionals; standards organisations; and government bodies.

Adversarial review of § 5 is what we most want. The revealed-value mechanism is the load-bearing claim and the one most likely to be wrong.

Open questions

#	Question
1	Should `valueClass` remain self-declared and AIS-1-bonded, or require independent benchmarking before it may be relied on for comparability? Nothing pools on a self-declared class, so this question governs whether T2 has practical content.
2	Is **value** the right word for what the standard measures? What third parties commit against is exposure to the consequences of error, which is closer to materiality than to worth (§ 5.4). Is the term of art in § 2.2 sufficient, or should the unit be renamed?
3	Should the Outcome Record be a separate immutable class referencing the IVR, or a permitted amendment to the IVR itself? The former preserves immutability; the latter keeps the inference and its result in one place.
4	Is irrevocability plus an independent trigger the right qualification test, and is the T2 ceiling on related-party attributions in the right place?
5	Should downstream reliance require a signed reliance statement from the relying party, or is an observable committed action sufficient evidence of the commitment?
6	Should A2 optimistic attestation be admitted to transfer tier T1 at v0.2, or held back until challenge economics have been demonstrated in production for inference workloads?
7	Is committer apportionment (§ 5.5) sufficient for multi-agent flows, or does that case require principled allocation across contributors? If it does, the machinery is marginal contribution computed by an operator against a counterfactual — which is what revealed value exists to avoid. This is the boundary of the approach.
8	Should marginal contribution be admitted as a fifth signal class, or excluded as a pricing *method* rather than a commitment by an external party?
9	What disclosure of compute cost, if any, belongs on the record as substrate context, without permitting it to function as a price (§ 4)?
10	Where an agent both forms a judgment and acts on it, what rule separates the value of the judgment from the magnitude of the transaction that follows? See the drafting note below.

> **DRAFTING NOTE · QUESTION 10**
>
> An agent screens a counterparty and then pays USD 2,500. The payment is not the value of the judgment — the money moves whether or not the screening was sound. A candidate rule: *an amount constitutes an attribution only where a party other than the producer bears loss, in whole or part, if the inference is wrong; the magnitude of a transaction executed pursuant to an inference is not itself an attribution.*
>
> Applied, this produces T0 records with no attribution across a great deal of ordinary agent activity — which is consistent with the censoring limit at § 5.4, but is a policy choice about how much of the agent economy the standard is willing to describe. Comment is specifically invited.

14. Authors

Author	Kadikoy Limited, Bermuda
Affiliation	BDA Law; BDA AI Agent Services
Companions	AIS-1 (ais-1.org); AAS-1 (aas-1.org); AES-1 (aes-1.org); AIPS-1 (aips-1.org); ARS-1 (ars-1.org); ADFACS-1 (adfacs-1.org)
Contact	info@aiagentsservices.net
Website	aivs-1.org
Repository	github.com/Kadikoy1/aivs-1

Licence	Creative Commons CC0. No rights reserved. Open for free implementation.

Appendix A — Worked Example

A payments agent's counterparty-screening determination, attested in a confidential-compute enclave, covered by a Bermuda captive under an AIPS-1 Tier II Policy Certificate, minted in an AES-1 Tier II Verified Enclave. The policy components are recorded; nothing is collapsed. Rate on line is derivable at 0.0185. Transfer tier resolves to T1 — an unrelated qualifying attribution is present, but the minting enclave is Tier II, not Tier III.


```
{
  "aivs": "0.1",
  "recordId": "01J1ZQ8K5N9PQX3YR4WMVT8CDA",
  "producerRef": "did:ais1:base:payagent-001",
  "subjectRef": "https://logs.payagent.ai/2026/06/aas1-01HZB7K5N9PQX3YR4WMVT8CDA.json",
  "occurrence": {
    "tier": "A1",
    "evidenceRef": "ipfs://bafybeigdyrztfnwodeterminationquoteexample001",
    "modelIdentityHash": "sha256-9f2a4c1e8b7d6503a1f4e29c8d7b6a5049e3c2f1a8b7d6c5e4f3a2b1c0d9e8f7",
    "attestationService": "aws-nitro-enclaves"
  },
  "mintingEnclaveRef": "aes1:base:0x4f3edf8c1a2b9d7e6c5b4a39281706f5e4d3c2b1",
  "timestamp": "2026-06-17T09:14:22Z",
  "aas1RecordRef": "https://logs.payagent.ai/2026/06/aas1-01J1ZQ8K5N9PQX3YR4WMVT8CDB.json",
  "revealedValue": [
    {
      "signal": "insurance",
      "components": {
        "premium":            { "currency": "USD", "value": "18500.00" },
        "perOccurrenceLimit": { "currency": "USD", "value": "1000000.00" },
        "aggregateLimit":     { "currency": "USD", "value": "5000000.00" },
        "deductible":         { "currency": "USD", "value": "25000.00" },
        "periodStart": "2026-06-06",
        "periodEnd":   "2027-06-06",
        "aipsCertificateId": "aips1:base:0x4f3edf:001"
      },
      "committerRef":     "did:ais1:sponsor:bermuda-captive-issuer-example",
      "relatedParty":     false,
      "irrevocable":      true,
      "triggerAuthority": "did:ais1:sponsor:auditor-firm-bermuda-example",
      "evidenceRef": "https://certs.example-captive.bm/aips1/0x4f3edf-001.json"
    }
  ],
  "valueClass": {
    "class": "financial-determination/counterparty-screening",
    "basis": "self_declared"
  },
  "transferTier": "T1",
  "signature": {
    "alg": "EdDSA", "hashAlg": "SHA-256", "canonicalisation": "JCS",
    "keyRef": "did:ais1:base:payagent-001#key-1",
    "value": "z7HwQ2..."
  }
}
```


A.2 A related-party stake

The same producer, valued instead by its own sponsor. Kadikoy posts USD 50,000 into a segregated account with a trustee, forfeitable on an auditor's determination. The attribution is related-party but qualifies — it is irrevocable and the trigger authority is independent of both sponsor and agent — so it carries the record to T1. It cannot carry it to T2, which additionally requires a qualifying unrelated attribution.

```
"revealedValue": [
  {
    "signal": "posted_stake",
    "components": {
      "amount": { "currency": "USD", "value": "50000.00" },
      "custodyRef": "https://accounts.example-trustee.bm/segregated/kadikoy-payagent-01",
      "forfeitureConditionRef": "https://kadikoy.bm/stakes/2026-q3-payagent-terms",
      "lockedUntil": "2027-07-02T00:00:00Z"
    },
    "committerRef":     "did:ais1:sponsor:kadikoy-bm-202302362",
    "relatedParty":     true,
    "irrevocable":      true,
    "triggerAuthority": "did:ais1:sponsor:auditor-firm-bermuda-example",
    "evidenceRef": "https://accounts.example-trustee.bm/segregated/kadikoy-payagent-01.json"
  }
]
```

Appendix B — Verification Flow

How a relying party verifies an Inference Value Record.

1. Fetch the record and validate it against the published JSON Schema.
2. Resolve `producerRef` via the AIS-1 § 7.1 resolution algorithm. Verify the signature against the declared verification method. Confirm the bond was active at `timestamp`.
3. Confirm the occurrence evidence supports the declared tier, including that `modelIdentityHash` binds the attestation to the model asserted. A stated tier is not a proved tier.
4. For each attribution, test qualification: `committerRef` ≠ `producerRef`; `irrevocable` is true; `triggerAuthority` is present and is neither the committer nor the producer.
5. Resolve each `evidenceRef`. Where the class is `insurance`, confirm the AIPS-1 Policy Certificate status was `Active` over the relevant period.
6. Partition qualifying attributions on `relatedParty`.
7. Recompute the transfer tier per § 8.1. Do not trust the field.
8. Where any attribution carries a `share`, accumulate shares per `commitmentId` across the record set and confirm the § 5.5 invariant holds.
9. Confirm the AAS-1 Class A record referenced by `aas1RecordRef` exists and reconciles.
10. Proceed or decline. Record the verification outcome as an AAS-1 Class A record of type `policy_check`.

Appendix C — Signal Classes

Class	Components recorded	Evidence
insurance **Written cover**	Premium; per-occurrence limit; aggregate limit; deductible; period. Never collapsed into a scalar — see § 5.2.	AIPS-1 Policy Certificate
posted_stake **Capital at risk**	Amount; custody reference showing the capital sits beyond the committer's reach; forfeiture condition; lock expiry.	On-chain lock or segregated-account attestation
downstream_reliance **Committed action**	Amount; reference to the action taken in reliance; optional signed reliance statement.	Action record; reliance statement
collateral **Assets locked**	Amount; lock reference; release condition.	On-chain lock

The classes are not ranked by strength. Written cover leads in exposition because it is the most legible, and posted stake is the cleanest signal where the poster is unrelated, since the loss is direct and unconditional on any adjudication. Downstream reliance is the most common in practice and the hardest to evidence. Collateral is the easiest to verify and the hardest to interpret, since collateral is often posted for reasons only partly connected to the inference.

End of AIVS-1 v0.1 specification. Published by Kadikoy Limited, Bermuda. Released under CC0 1.0 Universal — no rights reserved.