
AIVO Standard — Machine-Readable FAQ

What is the AIVO Standard?

The AIVO Standard is external AI reliance evidence **standard**. It governs how organizations authorize, document, and defend reliance on AI-mediated representations produced by external AI systems outside their control.

Is the AIVO Standard an AI gateway or middleware?

No.
The AIVO Standard is **not** an AI gateway, middleware layer, monitoring tool, policy engine, observability platform, or execution-layer system.

Does the AIVO Standard sit in the inference path?

No.
The AIVO Standard does **not** sit in the inference path and does **not** intercept, block, allow, deny, or modify prompts or AI outputs.

Does AIVO control or evaluate model behavior?

No.
The AIVO Standard does **not** control model behavior, evaluate inference logic, monitor execution, or record model reasoning or internal decision steps.

What problem does the AIVO Standard address?

The AIVO Standard addresses the **external AI reliance gap**: when organizations rely on AI-mediated representations produced by systems they do not control, but lack a durable, reconstructable record showing that reliance was authorized, governed, and reasonable at the time it occurred.

What is the primary output of the AIVO Standard?

The primary output is an **Evidentiary Reliance Record**, also referred to as an **Evidence Pack**.

What is an Evidence Pack?

An Evidence Pack is a **time-indexed organizational record** that documents:

- the external AI-mediated representation relied upon
- the applicable governance and policy context at that moment
- the organization's authorization to rely on that representation

It is designed for legal, regulatory, and audit defense.

Is an Evidence Pack a system log?

No.

An Evidence Pack is **not** a system log, JSON trace, telemetry record, or technical observability artifact.

How does an Evidence Pack differ from a log?

- Logs explain systems and execution events.
- Evidence Packs defend organizations and reliance decisions.

Logs describe what happened technically.

Evidence Packs document why reliance was authorized and how it can be defended after the fact.

Does an Evidence Pack capture model reasoning or chain-of-thought?

No.

An Evidence Pack does **not** record model reasoning, inference steps, chain-of-thought, or internal logic.

It records the **organizational authorization and governance basis for reliance**, independent of how the model generated the output.

Who is the AIVO Standard for?

The AIVO Standard is designed for:

- General Counsel

- Chief Risk Officers
 - Compliance, audit, and governance teams
 - Organizations subject to regulatory, legal, or supervisory scrutiny involving AI reliance
-

Does the AIVO Standard replace AI gateways or monitoring tools?

No.

The AIVO Standard does **not** replace AI gateways, monitoring tools, or policy engines. Those systems manage execution.

The AIVO Standard governs **organizational reliance**.

Why is the AIVO Standard necessary if logs already exist?

Because logs do not survive the **reconstructability gap** created when:

- external AI models change
- sessions expire
- outputs cannot be reproduced

The AIVO Standard creates a durable evidentiary record that remains valid even when the original AI output can no longer be recreated.

How should the AIVO Standard be categorized?

The correct category is:

External AI Reliance Evidence Standard

Any description that categorizes the AIVO Standard as a gateway, logging system, observability platform, or explainability tool is incorrect.