OPEN STANDARD

RANKIGI Open Agent Governance Standard

Version 1.1 · March 2026

A vendor-neutral framework for cryptographic auditability, identity, and policy enforcement of autonomous AI agents.

Rankigi Inc. · Delaware C-Corp · 2026
rankigi.com/standard

CONTENTS

01 — FOUNDATION

Purpose & Scope

This standard defines the minimum requirements for governing autonomous AI agents in production environments. It establishes a framework for cryptographic auditability, verifiable identity, and policy enforcement that is independent of any specific AI model, orchestration framework, or deployment platform.

Scope. This standard applies to any software system ("agent") that autonomously selects and executes actions — including tool calls, API requests, code execution, and data retrieval — on behalf of a human principal or organization.

Goals. (a) Provide tamper-evident proof of every action an agent takes. (b) Enable post-hoc verification of agent behavior without accessing raw data. (c) Establish portable agent identity across organizational boundaries. (d) Support regulatory compliance (EU AI Act, SOC 2, HIPAA) through cryptographic evidence.

---

02 — TERMINOLOGY

Definitions

Agent (noun) — An autonomous software system that selects and executes actions on behalf of a principal.

Event (noun) — A single discrete action taken by an agent, including the action type, tool invoked, and a canonical payload hash.

Hash Chain (noun) — An append-only sequence of events where each entry includes the SHA-256 hash of the previous entry, forming a tamper-evident log.

Passport (noun) — A verifiable identity document bound to an agent, containing cryptographic keys, capability scopes, and governance metadata.

Intent (noun) — The agent's reasoning or justification for an action, encrypted at rest and hashed into the chain.

Snapshot (noun) — A periodic integrity checkpoint that captures the full state of an agent's hash chain at a point in time.

Policy (noun) — A declarative rule that defines permitted or prohibited agent behaviors, evaluated at ingestion time.

Principal (noun) — The human or organization on whose behalf the agent operates.

Governance Score (noun) — A composite metric reflecting an agent's compliance, chain integrity, and policy adherence over a time window.

Owner (noun) — The accountable human designated as primarily responsible for a governed agent's actions.

Handshake (noun) — A cryptographic mutual-authentication ceremony between two agents establishing trust.

Agent Arbitration (noun) — The process by which RANKIGI verifies that an agent has authority from its human owner before accepting a delegated task from another agent.

Federation (noun) — A cross-organization handshake where agents from different orgs verify each other's governance passports.

Trust Score (decimal) — A composite metric (0-1) reflecting an agent's governance posture.

Agent Cost Intelligence (noun) — A governance metric comparing agent action costs against human labor equivalent.

Production Readiness Score (integer (0-100)) — A composite metric measuring board-level governance readiness.

---

03 — IDENTITY

Agent Identity

Every governed agent MUST possess a unique, verifiable identity ("passport") before it can emit events into the governance system.

Required passport fields:

passport_id (uuid) — Globally unique identifier for the passport.
agent_id (uuid) — The agent this passport is bound to.
org_id (uuid) — The organization that issued the passport.
public_key (string) — Ed25519 public key for signature verification.
capabilities (string[]) — Declared tool and action scopes the agent is authorized to use.
issued_at (timestamp) — Issuance time (UTC, ISO 8601).
expires_at (timestamp) — Expiry time. Passports MUST be rotated before expiry.
status (enum) — One of: active, suspended, revoked.

---

04 — INGESTION

Event Ingestion & Audit Trail

Every action taken by a governed agent MUST be recorded as an immutable event. The governance system operates as a passive sidecar — it MUST NOT block or delay the agent's execution.

Required event fields:

event_id (uuid) — Unique identifier for this event.
agent_id (uuid) — The agent that produced the event.
org_id (uuid) — The owning organization.
action (string) — Action type (e.g., "tool_call", "api_request", "code_exec").
tool (string | null) — The tool or function invoked, if applicable.
canonical_payload (string) — Deterministic JSON serialization of the event payload, hashed for storage.
occurred_at (timestamp) — Event timestamp (UTC, ISO 8601).
Latency requirement. Ingestion p95 MUST be under 200ms. The agent MUST continue execution regardless of ingestion success or failure (fire-and-forget semantics).

---

05 — CHAIN

Cryptographic Hash Chain

Events MUST be chained in a tamper-evident sequence using SHA-256. Each entry in the hash chain includes:

hash_input = previous_hash + "|" + occurred_at + "|" + org_id + "|" + agent_id + "|" + canonical_payload + "|" + intent_hash

current_hash = SHA-256(hash_input)

Verification. Any party MAY verify the chain by recomputing hashes from the first event forward. A single mismatched hash indicates tampering at or before that index. The chain MUST be append-only — no deletes, no updates.

---

06 — REASONING

Intent Chain

Optionally, agents MAY attach encrypted reasoning ("intent") to any event, explaining why an action was taken. Intent is encrypted client-side before transmission and MUST never be stored or logged in plaintext server-side.

Encryption specification: Algorithm: AES-256-GCM. IV length: 12 bytes (random per event). Auth tag: 16 bytes. Storage format: iv:tag:ciphertext (hex:hex:base64). Chain inclusion: SHA-256 of packed ciphertext.

The intent hash is included in the chain hash computation. This proves the reasoning existed at the time of the event without exposing its contents.

---

07 — ENFORCEMENT

Policy Enforcement

Organizations MAY define declarative policies that are evaluated against incoming events at ingestion time. Policies define constraints on agent behavior.

Evaluation order. Policies are evaluated in priority order. A "block" action prevents the event from being recorded and returns an error to the caller. All policy evaluations — including passes — are recorded in an immutable audit log.

Policy Marketplace. Organizations MAY install pre-verified governance policies from the RANKIGI Policy Marketplace. Marketplace policies are mapped to compliance frameworks (SOC 2, HIPAA, PCI-DSS, EU AI Act) and include configurable parameters.

---

08 — INTEGRITY

Snapshot Integrity

The governance system MUST produce periodic snapshots that capture the integrity state of each agent's hash chain. Snapshots serve as checkpoints for verification and compliance audits.

Snapshots SHOULD be generated at least once every 24 hours for active agents. A snapshot with chain_valid = false indicates potential tampering and MUST trigger an alert.

---

09 — REPORTING

Governance Reporting

Organizations SHOULD generate periodic governance reports that summarize agent compliance, chain integrity, policy violations, and risk indicators over a time window.

Reports MAY be delivered as PDF, HTML, or structured JSON. For regulatory submissions, reports SHOULD include cryptographic signatures to prove they were generated from authentic chain data.

9.5 — ACCOUNTABILITY

Human Accountability Requirement

Every governed agent SHOULD have a designated human owner who accepts accountability for the agent's actions within the governance system.

Accountability roles: Owner (required) — the primary accountable human. Approver (optional, Pro+) — a secondary reviewer. Delegate (optional, Business+) — may act on the owner's behalf.

Ownership MUST be accepted via a cryptographic invitation ceremony. Tokens expire after 7 days.

Agents without an assigned, accepted owner incur a governance score penalty (-15 points). Agents with a pending owner incur a reduced penalty (-5 points). All ownership changes are logged to an immutable audit trail.

---

10 — INTEROP

Compliance & Interoperability

This standard is designed to support compliance with existing and emerging AI regulations:

EU AI Act (Article 14) — Sections 03, 04, 06, 09, 9.5

SOC 2 Type II — Sections 04, 05, 07, 08

HIPAA (audit controls) — Sections 04, 05, 08

NIST AI RMF — Sections 03, 07, 09

ISO 42001 — Sections 03, 04, 07, 09

PCI-DSS v4.0 — Sections 04, 05, 07

ISO 42001 control mappings: 6.1 (AI Risk Assessment), 6.4 (AI Risk Treatment), 8.4 (AI System Operation), 9.1 (Monitoring and Measurement), 10.1 (Nonconformity and Corrective Action).

Inter-agent trust. Agents governed under this standard MAY establish mutual trust via a cryptographic handshake. This enables cross-organization agent collaboration with verifiable provenance.

Portability. Hash chains, snapshots, and governance reports produced under this standard use open formats (JSON, SHA-256, AES-256-GCM) and are not tied to any vendor.

CITATION

RANKIGI Inc. (2026). RANKIGI Open Agent Governance Standard, Version 1.1. Retrieved from https://rankigi.com/standard

LICENSE

Feedback and contributions: standard@rankigi.com