AI AGENT GOVERNANCE · GRC EDITION

Your GRC Framework Covers Human Users. Who Governs the AI Agents?

AI agents are running in your environments today — executing workflows, accessing sensitive data, calling external APIs — with no identity in your IAM, no controls mapped to your frameworks, and no audit evidence for your next assessment.

THE GOVERNANCE GAP

What Every GRC Team Is Missing Right Now

WITHOUT RUNTIMEAI

- ✕ AI agents have no identity in your IAM
- ✕ No controls mapped to SOC 2 / NIST / EU AI Act
- ✕ No audit evidence for agent actions
- ✕ No automated response when an agent goes rogue
- ✕ Non-Human User Risk domain is a blind spot
- ✕ Shadow AI invisible to compliance teams

WITH RUNTIMEAI

- ✓ Every agent has a verified, immutable identity
- ✓ All actions map to SOC 2, NIST AI RMF, EU AI Act controls
- ✓ Tamper-proof audit trail — Merkle-chained
- ✓ Kill switch fires in <50ms via dashboard
- ✓ Non-Human User Risk fully instrumented
- ✓ Shadow AI discovered automatically

THE PROBLEM IS ALREADY HERE

73%
of enterprises have AI agents running without governance controls

0
existing GRC platforms cover non-human identity governance natively

<10m
to deploy RuntimeAI and have full NHI visibility in your environment

HOW IT WORKS

🔍

Discover
Auto-detect every AI agent, MCP tool, and cloud workload — including shadow AI

🪪

Identify
Issue verified non-human identities with SPIFFE/X.509 certificates and risk scores

⚖️

Govern
Enforce policies via OPA. Map every action to your compliance framework controls

🔴

Respond
Kill switch fires in <50ms. Tamper-proof evidence packages ready for auditors

COMPLIANCE FRAMEWORK COVERAGE

Every AI Agent Action — Mapped to Your Controls

RuntimeAI auto-maps non-human agent behavior to the controls in your active compliance frameworks. No manual re-mapping. Evidence is generated continuously — not at audit time.

RuntimeAI Capability	SOC 2	NIST AI RMF	EU AI Act	Evidence Generated
NHI Discovery & Registry	CC6.1 Logical Access	GOVERN 1.1	Art. 9 Risk Mgmt	Identity issuance log
Policy Engine (OPA)	CC6.3 Authorization	MANAGE 2.2	Art. 14 Oversight	Policy decision log
Kill Switch	CC6.6 Incident Response	RESPOND 1.1	Art. 14 Human Oversight	Kill activation record
Immutable Audit Trail	CC7.3 Audit Logging	MAP 5.1	Art. 12 Record-Keeping	Merkle-chained events
Drift Detection	CC7.2 Monitoring	MEASURE 2.5	Art. 9 Continuous Mon.	Drift finding report
PII / DLP Shield	CC6.7 Data Protection	GOVERN 5.1	Art. 10 Data Governance	Redaction audit log

EVIDENCE ARCHITECTURE

Tamper-Proof Audit Trail

Every agent event — discovery, policy decision, risk escalation, kill switch activation — is captured in a **Merkle-chained immutable log**. Hash-linked blocks mean any tampering is immediately detectable. Auditors receive cryptographically signed evidence packages.

Framework-Aligned Evidence Packages

On demand or scheduled, RuntimeAI exports structured evidence packs aligned to your active frameworks — SOC 2, NIST AI RMF, EU AI Act, ISO 42001, FedRAMP, HIPAA. Import directly into your GRC platform or deliver to auditors as-is.

FRAMEWORKS SUPPORTED

SOC 2 Type I & II · NIST AI RMF 1.0 · EU AI Act · ISO 42001 · ISO 27001 · FedRAMP Moderate · HIPAA Security Rule · NIST CSF 2.0 · GDPR · PCI-DSS · CMMC 2.0 · + 49 more

BUILT FOR GRC & COMPLIANCE TEAMS

Three Ways RuntimeAI Fits Your Program

FOR GRC CONSULTANTS

Deploy for clients in under an hour. Export evidence packages mapped to each client's active frameworks. Runtime-tested controls — not just documented ones.

Your ACET-style tests now have real AI agent data to validate against.

FOR IN-HOUSE GRC TEAMS

Close the non-human identity gap in your IAM and GRC programs. Continuous compliance monitoring for every AI agent — not just point-in-time scans.

Evidence ready before your auditor asks for it.

FOR GRC PLATFORM BUILDERS

RuntimeAI is the AI agent data layer your platform is missing. API-first integration. We generate the non-human identity evidence — you surface it in your UI.

OEM and integration partnerships available.

DEPLOYMENT

Flexible Deployment Models

- ✓ **SaaS** — fully managed, up in <10 minutes, zero infrastructure
- ✓ **Private cloud** — runs in your Kubernetes cluster, your VPC
- ✓ **Air-gap / sovereign** — no external dependencies, full data residency control
- ✓ **Multi-tenant** — manage multiple client environments from one console

What You See in 10 Minutes

1. All AI agents discovered and ranked by risk score
2. Shadow AI surfaced — agents not in any IAM or CMDB
3. Compliance gap report against your active frameworks
4. Kill switch active — any agent can be stopped in <50ms

Ready to close the AI governance gap?

Start with a free discovery scan against your environment. See every AI agent, risk-scored and framework-mapped, in under 10 minutes.

runtimeai.io roshan@runtimeai.io calendly.com/roshan-runtimeai/30min