The Trust Crisis in Agentic AI

Five questions every CISO is asking that no current tool fully answers:

- Who approved this agent to access our claims database, and can I prove it to an auditor?
- Can I stop it right now, in under a second, if it deviates from its approved behavior?
- My business teams are deploying agents faster than my security team can review them. How do I scale governance without becoming the bottleneck?
- 150,000 agents per enterprise by 2028, each with a persistent credential to rotate, protect, and audit. Is there an alternative?
- Our identity, secrets, EDR, and compliance tools each solve a slice. Who converges them into a single governance decision before the first credential is ever issued?

Why Traditional IAM and Governance Can't Solve This

The enterprise security stack was built for humans logging into applications. Agentic AI breaks that model at the root. Incumbent vendors are retrofitting identity platforms designed for human access patterns onto autonomous agents operating at machine speed. The result: discovery-first governance that detects problems after deployment, persistent credentials that multiply the attack surface with every agent deployed (Gartner projects 150,000 per enterprise by 2028), and review models that can't keep pace with business adoption. No convergence across the six stakeholders who should govern an agent's lifecycle. No runtime enforcement tied back to what the business actually approved.

Monitoring tools will spot bad agent behavior after deployment. But without a governance baseline to compare against, they generate noise, not trust.

Runtime observability is necessary. It is not sufficient. It's a lagging indicator when the root of trust was never established.

Six Prerequisites for Provable Trust

TrustwAIre enters at agent design time, before a line of code runs, and enforces continuously from there.

Business Intent-Driven Governance. Proactive, not reactive. Business owners define agent purpose; CISOs enforce it. Governance flows from intent, not from after-the-fact log inspection.

Transparent Visibility & Decisioning. Business owner to CISO to SOC, connected in real time. Every approval, override, and exception visible across the chain. No silos, no surprises.

Cryptographic Enforcement. Every agent credential is cryptographically bound to a specific person, device, location, time window, and workflow. It cannot be reused, shared, or forged. When the operation completes, the credential ceases to exist.

One-Time Use Credentials. TTL and workflow-bound expiration. The credential dies, the access dies. No persistent tokens to steal. Blast radius contained to a single operation and time window.

Real-Time Chain of Custody. When something goes wrong, trace it to the exact moment, the exact actor, and the exact authorization state in seconds. Compliance evidence is a byproduct of normal operations, not a forensic reconstruction.

Immutable Auditability. The audit record cannot be altered by anyone, including TrustwAIre. Any tampering is mathematically detectable. What regulators see is exactly what happened.

How It Works

The TrustwAIre Credential Router™ coordinates OPA policy enforcement, SPIFFE/SPIRE X.509 certificate issuance, and iVALT biometric human attestation in a single runtime transaction. Six stakeholder roles converge on a living governance profile before the first credential is issued. The CISO publishes thresholds; business owners operate within bounds. Agents receive ephemeral, resource-scoped PKI credentials with no token to rotate, no central callback, and no persistent access to secure.

A behavioral Kill Switch compares live agent actions against the governance baseline and terminates in under one second with cryptographic proof. Compliance evidence is generated as a byproduct of doing the governance, not collected after the fact.

Built on vendor-neutral CNCF open standards (SPIRE, OPA, Envoy, Falco). Integrates with your existing identity, secrets, EDR, and SIEM infrastructure. Adds the governance layer none of them provide.

Next Steps

See TrustwAIre in action: **[20-Minute Video Walkthrough + Multi-Part Demo Series]** · **[Download the Evaluation Build]** · **[Read the Position Papers and related collateral]**

TrustwAIre · Secure AI LLC · contact@trustwaire.ai · trustwaire.ai