Tumeryk
AI Trust Score™

www.tumeryk.com

AGENTIC AI GOVERNANCE

Governance and Security for Autonomous AI systems

DISCOVER | **ASSESS** | **SECURE** | **GOVERN**

Agentic AI introduces a new class of enterprise risk. Unlike traditional software, autonomous AI systems can independently access data, invoke APIs, and take actions across business systems. Without governance and cybersecurity controls, Agentic AI can expose sensitive data, violate regulations, propagate errors at scale, and cause material harm to the business.

Tumeryk provides a real-time trust, governance, and security layer for agentic AI, powered by the AI Trust Score™. Every agent action, decision, and outcome is continuously assessed for security, compliance, and policy risk, giving enterprises clear visibility and control over autonomous AI behavior while protecting enterprise data and critical systems.

Tumeryk enables consistent governance across AI agents, copilots, and orchestration workflows through embedded policy enforcement, behavioral monitoring, and continuous oversight. This ensures agentic AI operates safely within defined enterprise risk boundaries, allowing organizations to innovate with autonomy without compromising data security, compliance, or business integrity.

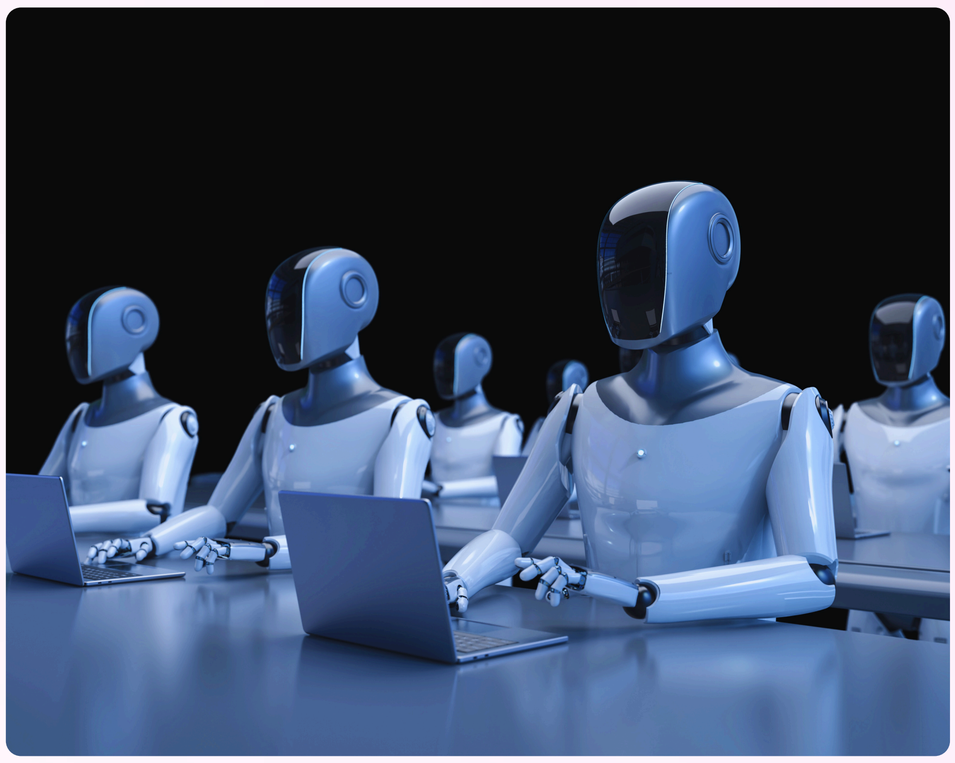

AI TRUST = AI SECURITY + AI SAFETY + RESPONSIBLE AI

Agent Access Control: Policy-based control over agent permissions, actions, and execution paths to ensure autonomous AI systems operate within approved boundaries.

Observability: Real-time inspection of agent actions, decisions, and outcomes to detect risky, non-compliant, or off-policy behavior across AI workloads.

Runtime Guardrails: Automated enforcement of policies to prevent unsafe actions, data misuse, or unintended behaviour during agent execution.

Model Risk Audit: Ongoing red-teaming, testing, and audit-ready visibility into Foundational AI Models to support governance, compliance, and reporting.

Tumeryk AI Trust Score™ Platform

Tumeryk is the trust infrastructure for AI. Tumeryk delivers enterprise-grade security and compliance for agentic and conversational AI. Its AI Trust Score™ and policy engine help organizations monitor, block, and manage AI behavior in real-time, ensuring accuracy, safety, and compliance.

KEY FEATURES

- Redteaming for AI Models
- Multi-cloud security for agentic AI workloads
- Real-time scoring of agent actions and outcomes
- Centralized oversight for enterprise AI workloads
- Easy integration with existing AI and cloud platforms
- Detection of hallucinations and unsafe agent outputs

BENEFITS

- Ensure continuous agent governance
- Compliance and vulnerability detection
- Enforce accountable AI decision making
- Block unauthorized agent execution paths
- Prevent unsafe or unintended agent actions
- Establish trust in agentic AI systems at scale

sales@tumeryk.com

aws Available in AWS Marketplace

GTM PARTNER

pwc

NVIDIA INCEPTION PROGRAM

WIZ CLOUD WIN CERTIFIED INTEGRATION

Gartner COOL VENDOR 2025