RED VECTOR

PRODUCT SHEET

AGENTWATCH™

Your AI agents are acting at machine speed. **AGENTWATCH™ governs them before they govern you.**

METHODOLOGY
IRM applied to AI entities

ALIGNMENT
OWASP Top 10 Agentic AI

TIME TO VALUE
30-day Bootcamp

CATEGORY
AI Agent Risk Governance

AGENTWATCH™ applies proven Insider Risk Management principles to AI entities, so security leaders can detect, investigate, and respond to agent risk before it becomes a breach. Every agent is treated as a **Digital Insider**: an autonomous synthetic employee that must be continuously monitored, baselined, and governed.

THREE BLIND SPOTS IN YOUR AI DEPLOYMENT

THE IDENTITY GAP

Agents operate as service accounts or API identities without ownership tracking or lifecycle governance.

THE BEHAVIOR GAP

Agents accumulate privileges, expand scope, and evolve. Traditional tools cannot baseline these risk patterns.

THE GOVERNANCE GAP

No unified way to monitor AI activity, enforce policy, and link agent actions to accountable humans.

The more autonomous AI you deploy, the harder it becomes to know **who, or what, is acting on your behalf.**

FEATURES

Agent Discovery

Inventories sanctioned and unsanctioned AI agents.

Behavioral Analytics

Per-agent baselines that flag activity drifting outside its normal pattern.

Risk Scoring

Composite scoring across AI Agent Risk Indicator domains.

Owner Attribution

Links agent actions to the accountable humans.

Analyst Workbench

Single product for triage, inquiry, and investigation for AI Risk Analysts.

Automated Response

Acts on AI agents of concern when risk reaches an unacceptable level.

WHY IT'S DIFFERENT

- **IRM applied to AI.** Proven Insider Risk Management methodology adapted to govern non-human AI entities.
- **AI Risk Indicators (ARIs).** Composite scoring across Data Security, Identity, AI Behavior, Governance, Threat Exploitation, and Human-AI Alignment.
- **Stack integration, not replacement.** Leverages your existing SIEM, IAM, DLP, and AI observability investments.
- **Continuous behavioral baselining.** Detects drift as it happens, not after the fact.
- **Human-AI accountability loop.** Full traceability from agent action to the employee who deployed it.
- **AI Agent Risk Maturity Model.** A structured path from ad-hoc visibility to board-level predictive AI risk reporting.

PROOF POINT

"We had no way to know which AI agents were acting inside or outside their sanctioned scope. AGENTWATCH™ gave us a single risk picture across our entire AI deployment and the ability to act before a breach occurs."

— Fortune 500 Financial Services Enterprise

WHO IS THIS FOR

- **CISOs** combating shadow AI
- **CROs / CCOs**: EU AI Act, GDPR, NIST AI RMF, ISO 42001
- **Insider Risk Analysts**: human intent × AI execution
- **SOC Analysts**: real-time detection, agent isolation
- **CTOs**: faster, safer AI deployment
- **GRC & AI Ethics Teams**: model cards, bias controls, compliance

BEGIN YOUR ENGAGEMENT

The Red Vector Bootcamp applies AGENTWATCH™ to your environment in 30 days: validated risk report, business case, and action plan delivered to your CISO.

Start with a Bootcamp

redvector.ai · info@redvector.ai

© 2026 Red Vector, Inc. · info@redvector.ai · redvector.ai