



The AI Agent Blast-Radius Assessment

Find out exactly what your AI agents can do, who authorized it, and what breaks if one is compromised. A fixed-fee, senior-led assessment that inventories your agent estate, maps its blast radius against the 12 Principles for Agent Runtime Governance, and hands you a prioritized roadmap and a board-ready report. In weeks, not quarters.

What We Assess

Agent & MCP Inventory

Every agent and MCP server in scope, including shadow AI, with owner and purpose.

Delegation & Authority Chains

How authority propagates across agents and where it silently widens.

Auditability & Evidence

Whether you can reconstruct and prove what any agent did.

Standing Privilege & Blast Radius

What each agent's credentials can reach, and the damage one compromise could do.

Runtime Authorization Gaps

Every control measured against the 12 Principles for Agent Runtime Governance.

Maturity Score & Roadmap

A 0–3 maturity score and a prioritized path to close the gaps.

What You Walk Away With

- ✓ Complete agent and MCP inventory with ownership
- ✓ Standing-privilege and blast-radius map per agent
- ✓ Delegation and authority-expansion analysis
- ✓ Maturity score against the 12 Principles (0–3)
- ✓ Prioritized 30/60/90 remediation roadmap
- ✓ Board-ready executive report and technical plan

Engagement

- ✓ Week 1 — Discovery & inventory
- ✓ Weeks 2–3 — Analysis & blast-radius mapping
- ✓ Weeks 3–4 — Executive readout & roadmap

Inside the Executive Report

Board-ready deliverable

- 01 Executive summary and risk posture
- 02 Agent inventory and shadow-AI findings
- 03 Standing-privilege and blast-radius exposure
- 04 Delegation and authority-expansion risks
- 05 Auditability and evidence gaps
- 06 Maturity score and peer benchmark
- 07 Prioritized 30/60/90 remediation roadmap
- 08 Appendix: methodology and evidence

Fixed Scope, Fixed Price

Rapid Scan

from \$15,000

1 week

One team or framework. A fast read on your biggest exposure.

Standard Assessment

from \$40,000

2–3 weeks

Organization-wide. The full picture, a maturity score, and a board-ready plan.

Enterprise

Custom

4+ weeks

Multi-BU and framework, a compromised-agent tabletop, and a path to Beacon.

Pricing is indicative and scoped per environment. The assessment credits toward a Watchlight AI Beacon design-partner engagement.

We wrote the framework

Authors of the 12 Principles for Agent Runtime Governance.

23+ years in the trenches

Enterprise identity and security for the largest, most regulated organizations.

A method, not an opinion

The same maturity model and controls that define the discipline.

Know your blast radius before an attacker does.

Start with a 30-minute scoping call. We size the engagement and confirm the fee.

watchlight.ai/consultation