SHADOW AGENT DISCOVERY

The agent blind spot

Most organisations run more AI agents than they can currently name. Copilots, workflow bots, RPA-plus-AI, agents a team switched on without telling anyone else. Almost nobody can currently produce a full, accurate list of either.

91%

of organisations run AI agents today. Only **10%** have a well-developed strategy to manage them.

Okta, AI Agents at Work, 2026

82%

discover an AI agent or workflow that IT and security did not know about.

Cloud Security Alliance, 2026

Why this is a blind spot, not just a backlog item

An agent decides for itself what to do next, so you cannot list every action it might take in advance. It can be misled by the content it reads. Its permissions can widen over time without a code review. And there is no way to sign off its behaviour once, the way you might approve a fixed piece of software. That is why an inventory built the old way, by asking teams what they have deployed, reliably undercounts what is actually running.

Three questions worth asking before anyone else does

1. Do you have an AI strategy?

A named owner for agent risk, a policy for what is allowed, and a clear line for what needs human approval.

2. Do you know what is out there?

A full inventory, including the agents nobody logged, and which act for a person versus run on their own.

3. Is any of it governed?

Whether your identity foundation is solid enough to build on, and whether an agent can ever see more than the person behind it.

What we can honestly say: we do not claim to secure every agent in your estate, and you should be wary of anyone who does. What we can do is help you discover, control access to, and govern your agents, your non-human identities, and your secrets, using what is already proven and deployable today, and tell you honestly where that coverage stops.

Find out where you stand

Free, 30 minute Agentic Identity Readiness Assessment. No product pitch, no fixed scope.

Book the readiness assessment

intragen.com/what-we-do/ai-identity-risk