Ekselens

EKSELENS360

EKSELENS INSIGHT · AI GOVERNANCE

Agentic AI governance: every agent needs an owner, an identity and an off-switch

AI agents act, not just advise. Govern them like a workforce: a register, a named owner, a scoped identity, escalation thresholds and a tested off-switch.

THE EIGHT-PHASE METHOD

01	02	03	04	05	06	07	08
Discover	Diagnose	Design	Plan	Mobilise	Deliver	Embed	Optimise

Ekselens Consulting

ekselens.com · info@ekselens.com

Changing how transformation is done

The EKSELENS360 family

AI GOVERNANCE · 12 NOVEMBER 2026 · 4 MIN READ

AI has moved from answering questions to taking actions. Agents now raise purchase orders, triage incidents, update records and draft customer responses with limited human involvement. The governance most organisations built for generative AI — an acceptable-use policy and a list of approved tools — was designed for advice. It does not hold when software acts on the organisation's behalf.

The question boards should now ask is simple: for every agent running in our estate, can we say who owns it, what it is allowed to do, and how we would stop it?

Why agents change the risk picture

A chatbot that gives a poor answer creates a quality problem. An agent that takes a poor action creates an operational, financial or regulatory event. Three things change when AI becomes agentic:

- **Actions compound.** An agent can execute hundreds of steps before a human looks at the output. Errors propagate faster than review cycles.
- **Accountability blurs.** When an agent acts through a shared service account, the audit trail shows the system, not the decision-maker.
- **Scope drifts.** Agents are often granted broad permissions during a pilot "to make it work" and keep them in production.

None of this is a reason to slow down. It is a reason to put the controls in the build, not in a policy document that nobody reads at deployment time.

Start with an agent register

You cannot govern what you cannot count. The first control is an inventory of every agent in use — built in-house, embedded in a SaaS platform or switched on by a supplier. For each agent, record:

- the business process it supports and the decisions it can take;
- the named business owner accountable for its outcomes;
- the systems and data it can read and write;
- its risk tier, based on the impact of a wrong action;
- the date of its last review.

Most organisations discover agents they did not know they had, usually inside platforms they already licence. The register is also the foundation for regulatory work: the EU AI

Act's high-risk obligations now apply from 2 December 2027 for stand-alone systems and 2 August 2028 for AI embedded in regulated products, following the Digital Omnibus amendments. The extra runway is time to build evidence, not a reason to defer it.

Give every agent its own identity

Shared credentials are the single most common weakness in agent deployments. Each agent should have a distinct, scoped identity, with permissions limited to the task, credentials that expire and rotate, and every action logged against that identity. This is ordinary identity and access management discipline applied to a new kind of worker. It makes attribution possible, and attribution is what turns an incident into a fixable problem rather than an argument.

Design escalation and the off-switch before go-live

Every agent needs two things defined before it goes live:

- **Escalation thresholds.** The point at which the agent must stop and hand to a human — a value limit, a confidence level, a category of customer, or an action that cannot be reversed.
- **A tested off-switch.** A documented way to suspend the agent quickly, with a fallback process the business can run manually. "We could turn it off" is not the same as having tested it.

Tie both to the risk tier in the register. A low-risk drafting agent may need light-touch sampling. An agent that commits spend or changes customer records needs hard limits and real-time monitoring.

Make assurance continuous, not a one-off approval

Agents change. Models are updated, prompts are edited, tools are added and data drifts. A single approval at launch decays quickly. Replace it with a living cycle:

- review each agent against its register entry on a fixed cadence set by risk tier;
- sample decisions and compare them with what a competent human would have done;
- track incidents, overrides and escalations as leading indicators;
- retire agents formally, revoking identities and access, rather than leaving them dormant.

This is where the System of Proof matters. If every agent decision is traceable to an owner, a permission and a review, the organisation can show its board, its auditors and its regulator that control is real rather than asserted.

What leadership should commit to now

For most organisations, a credible first 90 days looks like this:

- **30 days:** build the agent register and name an owner for every entry.
- **60 days:** move every agent onto its own identity and remove shared service accounts.
- **90 days:** define escalation thresholds and test the off-switch for every high-impact agent.

None of these steps needs a new platform or a large programme. They need clear ownership and the discipline to apply existing controls to a new kind of actor.

Agentic AI will create real productivity gains. The organisations that keep them will be the ones that can prove, at any moment, who is accountable for what their agents do.

Want to talk this through?

Book a 45-minute discovery conversation. Bring one challenge, leave with three prioritised next steps — written up within 24 hours.

ekselens.com/contact
info@ekselens.com
+44 (0) 20 7874 7217

© Ekselens Consulting. All rights reserved.