

The scale of invisible AI agents

Findings from the 2026 State of Agent Security Report reveal a significant governance challenge for regulated institutions. In the environments studied, approximately 1,280 third-party products now embed AI agents. Of these, only 282 authenticate through **single sign-on** systems, leaving roughly 1,000 AI agents operating outside the visibility of identity infrastructure.

This gap does not result from deliberate concealment. Rather, it reflects a fundamental limitation: identity stacks can only govern what authenticates through them, and most embedded AI agents never do. The agents operate within third-party products that may already have access to enterprise systems, but the AI components themselves remain unregistered with central identity controls.

Implications for Saudi institutions

For organizations subject to NCA and SAMA frameworks, this presents a control gap. **Third-party AI agents** that bypass identity infrastructure cannot be inventoried, monitored for privilege escalation, or audited through conventional identity and access management tools. The agents may process sensitive data, make decisions, or interact with other systems without appearing in access logs or governance dashboards.

The challenge is architectural. Security controls designed for AI systems an organization explicitly deploys do not extend to agents embedded in software already in use. As AI capabilities become standard features in enterprise software, the number of ungoverned agents will grow unless institutions adapt their identity and access frameworks to detect and manage agents that do not authenticate in traditional ways.

What to do now

1. Conduct an inventory of all third-party software and SaaS platforms currently in use to identify which products may embed AI agents
2. Request vendor documentation on AI capabilities within existing enterprise software, including authentication methods and data access patterns
3. Evaluate whether current identity and access management tools can detect or log activity from embedded AI agents that do not authenticate through SSO
4. Review contracts and service agreements with third-party vendors to clarify governance rights over embedded AI components
5. Establish a cross-functional working group involving IAM, procurement, and risk teams to develop detection and governance controls for non-authenticating AI agents
6. Update vendor risk assessment questionnaires to include specific questions about embedded AI agents and their integration with enterprise identity infrastructure