

Sponsored content | White paper

The CIO agentic AI governance checklist

Assess your organization's readiness to govern agentic AI and other non-human identities.



CIO

Sponsored by



Non-human identity governance has become a board-level priority. If a security, privacy, or compliance issue occurs, CIOs must be ready to explain — at a moment's notice — how decisions made by AI agents, machine accounts, or other non-human entities are governed and who is accountable for them.

Use the following checklist to evaluate your organization's AI agent governance readiness and identify gaps before they become business or audit issues. Each section includes key capabilities to verify and a practical example of what effective AI governance looks like in practice.

□ Govern non-human identities quickly.

- Identify AI agents and other non-human identities operating across your environment.
- Ensure there's a human responsible for every non-human identity, and that each is governed throughout their lifecycle, from creation through retirement.
- Apply the same governance standards to non-human identities that you apply to human identities.
- Evaluate whether you have the capabilities to govern AI agents at the speed at which they operate.

Example: An engineering team first deployed a "VendorPay-Bot" to automate invoice processing, and the enterprise's identity platform immediately detected it. Instead of letting the AI agent operate in the shadows, the system assigned the bot a unique non-human identity and placed it under a strict, default "sandbox" policy, helping ensure it was governed from the second it was created.

□ Establish clear human accountability.

- Assign a human owner to non-human identities.
- Maintain ownership as non-human identities evolve or responsibilities change.
- Trace AI decisions back to the person responsible for that identity.

Example: Before the VendorPay-Bot accessed live financial systems, the governance platform required a human owner. The bot was assigned to James, the director of procurement. If it requests new access or makes an error, the system knows who's accountable so they can rectify the situation right away.

□ Apply policies consistently to AI decisions.

- Define when AI can act independently and when human approval is required.
- Apply policies consistently to every AI-driven identity decision.
- Continuously validate access against risk tolerance to ensure minimal risk.

Example: The VendorPay-Bot routinely processes low-dollar invoices autonomously. When it encountered a vendor contract that exceeded the set limit of \$150,000, policy determined that the risk exceeded its authority. The system paused the bot and routed to James for approval.

❑ Provide evidence that AI governance is working.

- Explain AI access decisions.
- Identify policies that governed the decision.
- Produce audit-ready evidence showing who approved, validated, or owned the decision and which policies were considered.

Example: During the annual compliance review, an auditor flagged that \$150,000 transaction. James immediately pulled a “decision receipt” showing the bot’s identity, the exact policy triggered, the timestamp, and his explicit digital sign-off. This transformed a black-box AI action into verifiable, audit-ready evidence.

❑ Establish a plan to contain compromised AI agents.

- Demonstrate the process to sequester, shut down, or otherwise stop an agent that is hallucinating, has been hijacked by an outside party, or has gone rogue.
- Consider options such as quarantine, ending the session, or hitting the “kill switch” if needed.

Example: A configuration error caused the VendorPay-Bot to suddenly attempt to query the highly confidential HR payroll database at 3 a.m. Continuous monitoring detected this behavioral drift and triggered the kill switch. The system automatically revoked the bot’s active session tokens and disabled its directory account, helping contain a potential internal breach.

❑ Maintain continuous visibility and governance.

- Continuously discover and govern new, non-human identities across your environment.
- Detect and respond to changes in behavior or access that increase risk.
- Identify capability gaps that require new skills, new processes, or external expertise.
- Evaluate how you’re using AI-powered identity governance capabilities to keep up with the velocity of agentic decision-making.

Example: After the CIO demonstrated that the team successfully discovered, governed, audited, and safely quarantined the VendorPay-Bot throughout its life cycle, the board approved 50 new, automated supply-chain agents. The team successfully put in place the AI-powered governance capabilities required to expand the AI workforce safely by ensuring that these new AI agents will be governed properly throughout the day.

Put your assessment to work

Strong agentic governance starts with understanding your readiness. Identifying governance gaps early helps organizations deploy AI safely and scale with confidence.

[Explore how explainability, audit-ready context, and continuous governance](#) help organizations deploy agentic AI with confidence.

DISCLAIMER: THE INFORMATION CONTAINED IN THIS DOCUMENT IS FOR INFORMATIONAL PURPOSES ONLY, AND NOTHING CONVEYED IN THIS DOCUMENT IS INTENDED TO CONSTITUTE ANY FORM OF LEGAL ADVICE. SAILPOINT CANNOT GIVE SUCH ADVICE AND RECOMMENDS THAT YOU CONTACT LEGAL COUNSEL REGARDING APPLICABLE LEGAL ISSUES.