Everyone Wants an AI Agent. Few Can Govern One.

Use the checklist to decide what can scale

Assess each deployment against the work it performs and the consequences of its actions. An agent that reads a limited set of customer records needs different safeguards from one that can send those records externally or submit orders. Readiness depends on evidence that the controls required for that deployment work.

The checklist below brings the preceding questions together under Discover, Protect, and Manage. For each requirement, record a status, supporting evidence, an accountable owner, and the next action. Use Verified when evidence demonstrates that the requirement is met, Partially addressed when coverage is incomplete, and Unresolved when the requirement remains unaddressed or unverified.

Record the agent's name, approved purpose, assessor, and assessment date alongside the checklist so it can also serve as a standalone review document.

	Governance requirement	Status	Evidence	Owner	Next action
Discover	Identify the agent and distinguish active, experimental, duplicate, or abandoned deployments.				
	Map accounts, credentials, delegated permissions, accessible resources, and connections to other agents.				
	Detect new deployments and changes to existing access or connections.				
Protect	Assign current business and technical owners and define the agent's approved purpose and prohibited actions.				
	Limit permissions by resource, action, and duration, including elevated and delegated access.				
	Enforce policy before execution and require human approval for defined consequential actions.				
	Attribute actions to the agent and initiating user where applicable, retain decision records, and support suspension.				
Manage	Require registration and approval before granting enterprise access.				
	Reassess changes to purpose, tools, data access, delegation, ownership, or autonomy; identify inactive or unowned agents.				
	Verify that retirement revokes access and stops scheduled or delegated activity while preserving required records.				

Use gaps identified through the checklist to set deployment boundaries. Establish clear ownership before granting further access, and remove permissions that exceed the approved task. Where runtime controls cannot enforce limits on consequential actions, restrict those capabilities or require human approval through an enforced workflow until the gap is closed.

Avoid reducing the assessment to a single readiness score. Several verified requirements cannot compensate for an uncontrolled ability to export sensitive data. Prioritize gaps according to the agent's access and potential impact, then use the evidence to decide whether to expand its scope, keep it within current limits, or pause affected capabilities.

Explore Saviynt Zuma

Saviynt