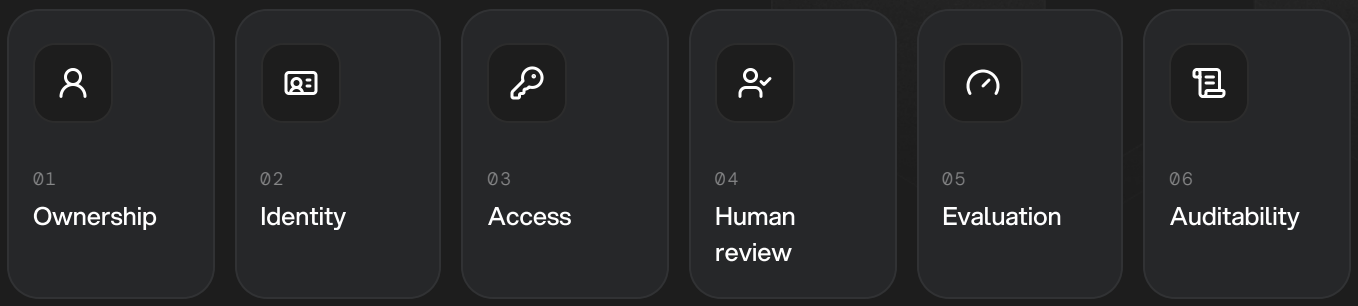
StackAI

GUIDE · 2026

28-POINT AGENT GOVERNANCE CHECKLIST

The CIO's guide to governing AI agents

Score your agent program across six key areas and see how ready you are for regulated, production-grade agents.



01 · INTRODUCTION

Agents are a new kind of identity. They deserve the same discipline as any critical system.

This checklist surfaces governance gaps before an auditor or a security review does, or before a real incident exposes them, such as an agent sharing confidential data or taking an action no one approved.

The market is reaching the same conclusion. On October 27, 2026, IBM announced identity and control capabilities for AI agents: a governed identity record per agent, short-lived least-privilege credentials, and one place to discover, contain and evaluate them.

150,000+

agents are expected to be deployed at the average Fortune 500 company by 2028.

SOURCE: GARTNER, APRIL 2026

Only 13%

of organizations say they have the governance to manage them.

SOURCE: GARTNER, APRIL 2026

HOW TO USE THIS GUIDE

STEP 01

Work through six areas

Ownership, identity, access, human review, evaluation and auditability.

STEP 02

Tick what is true today

Check a box only when the control is in place for every agent in production.

STEP 03

Score and act

Count your boxes out of 28 to see where you stand and what to fix first.

The takeaway. A governed agent has an owner, its own identity, only the access it needs, a human at the moments that matter, a tested accuracy target, and a complete audit trail.

01

Ownership

/ 5

Every agent has a named human who answers for it.

- [ ] Each agent has a business owner and a technical owner on record
- [ ] Ownership transfers when someone changes roles or leaves
- [ ] Each agent has a written purpose: the one job it is approved to do
- [ ] There is a full inventory of agents in production, including ones built outside your main platform
- [ ] Retired agents are decommissioned, not left running

02

Identity

/ 4

Agents are treated as non-human identities, not as extensions of a user or a shared service account.

- [ ] Each agent has its own entry in the company's identity system (such as Okta or Microsoft Entra), with its own credentials rather than the login or API keys of the person who built it
- [ ] No agent runs on a shared or personal credential
- [ ] Lifecycle status (draft, testing, production, retired) is tracked per agent
- [ ] Each agent's identity entry lists who owns it, what job it does and exactly what it is allowed to access or change

03

Access permissions

/ 5

Agents get only the access their job needs, for only as long as they need it.

- [ ] Each agent can reach only the data and systems its job requires, and nothing more
- [ ] Credentials are short-lived or rotated, never hardcoded in prompts or workflows
- [ ] Agents answer only from data the requesting user is already cleared to see
- [ ] Access can be revoked immediately
- [ ] Tools, MCP servers and integrations each agent can call are listed and approved

04

Human review points

/ 4

People approve the actions that matter before they happen.

- [ ] Consequential actions (sending, paying, writing to a system of record) require human approval
- [ ] Low-confidence outputs are routed to a human reviewer
- [ ] Reviewers know what they are approving and can reject or edit
- [ ] Escalation paths are defined for exceptions the agent cannot handle

05

Evaluation

/ 5

Agents are tested before launch and monitored after it.

- [ ] Each agent has a written accuracy target (for example, 95% correct on a test set of real cases) and must meet it before launch
- [ ] Test sets reflect real inputs, including edge cases and bad data
- [ ] Production accuracy is checked against that target on a set cadence, with an alert if it drops
- [ ] Model and token costs are tracked per agent
- [ ] Changes to prompts, models or tools trigger re-testing

06

Auditability

/ 5

You can reconstruct what an agent did, why, and on whose authority.

- [ ] Every run is logged: inputs, outputs, tools called and data sources used
- [ ] Answers cite their sources so reviewers can check them
- [ ] Writes to business systems carry provenance back to the agent run
- [ ] Approval decisions are recorded with who approved and when
- [ ] Logs are retained to match your regulatory requirements

02 · SCORE YOUR PROGRAM

Calculate your governance score

YOUR SCORE _____ / 28

STEP 1 · YOUR SCORE BY SECTION

SECTION	CHECKS	YOUR SCORE
Ownership	5	_____ / 5
Identity	4	_____ / 4
Access permissions	5	_____ / 5
Human review points	4	_____ / 4
Evaluation	5	_____ / 5
Auditability	5	_____ / 5
Total	28	_____ / 28

→ **Your weakest section:** _____ Start here, even if your total score is high.

STEP 2 · WHERE YOU STAND

0–13

UNGOVERNED

Agents ahead of controls

Agents are running without the basics.

→ Pause new launches. Start with ownership and access.

14–22

PARTLY GOVERNED

Gaps in one or two areas

Some controls exist, but not for every agent.

→ Fix the weakest section before adding agents.

23–28

GOVERNED

Ready to scale agents

Controls are in place across all six areas.

→ Expand to new teams and higher-risk workflows.

03 · VENDOR REVIEW

Questions to ask any agent platform provider

Use these in your next vendor review. A governed platform answers each one with a feature you can see, not a roadmap slide.

AREA	ASK
Ownership	Can I see every agent in production, who owns it and what stage it is in?
Identity	Do agents and users sign in through our SSO, with roles we control?
Access	Can an agent ever return data the requesting user is not cleared to see?
Human review	Where can I require approval before an agent acts, and who can approve?
Evaluation	How do I test an agent before launch and roll it out to a small group first?
Auditability	Can I trace a single answer back to every prompt, tool call and source behind it?

HOW TO RUN THE REVIEW

TIP 01

Ask to see it live

Every answer should be a working feature in the product, not a slide or a roadmap date.

TIP 02

Use your own workflow

Bring one real process, such as an IT support desk or KYC review, and walk it end to end.

TIP 03

Score them on the 28

Run the vendor through the same checklist you used to score your own program.

04 · PUTTING IT INTO PRACTICE

StackAI: from checklist to controls

Every area of the checklist should map to a control you can see in the product. Here is how StackAI, the enterprise platform for building, deploying and governing AI agents, handles each one.

CHECKLIST AREA	HOW STACKAI SUPPORTS IT
Ownership	Agent development lifecycle (ADLC) moves each agent from draft to production with a clear owner at every stage
Identity	SSO and role-based access control (RBAC) tie every builder, reviewer and user to a governed identity
Access permissions	RBAC, usage limits and permission-scoped answers that respect the access users already have across 500+ enterprise integrations. Deploy in our cloud, your cloud or on-prem
Human review points	Approval gates before consequential actions; low-confidence outputs flagged for review
Evaluation	Feature flags for controlled rollouts, plus analytics and monitoring on every agent in production
Auditability	Traceability down to each tool call and prompt, cited answers, and write-back with provenance

THE STACKAI PLATFORM

A powerful agent harness

Parses documents, websites, images and slides; orchestrates APIs, code and MCP servers; searches data with SQL, RAG and keywords; navigates visual interfaces with computer use.

Easy to build, end to end

No code. Go from idea to workflow to an exportable interface, with governance throughout. Forward deployed engineers help you find, prioritize, build and scale use cases.

Enterprise-grade by default

ADLC, RBAC, deploy anywhere, feature flags, analytics and monitoring, usage limits, and traceability to each tool call and prompt.

05 · STACKAI

Governed agents already at work across the enterprise

IT service management

IT support desk, account onboarding, PTO review

HR

Expense review, candidate screening, policy onboarding

Legal and compliance

KYC research and account approval, AML alert review, call center QA, engineering design QA, trade desk order review

Operations

Staffing, loan and claim underwriting, procurement copilots

Revenue operations

Sales support, RFP responses

Customer case study

A leading engineering and architecture firm runs staffing, legal review, engineering QA and HR agents on StackAI, with SSO, RBAC and monitoring across every one.

06 · GLOSSARY

Terms used in this guide

TERM	MEANING
Non-human identity	An identity for software (an app, service or agent) rather than a person
Least privilege	Granting only the access a task needs, and nothing more
RBAC	Role-based access control: permissions set by role, not by individual
SSO	Single sign-on through your company's identity provider
ADLC	Agent development lifecycle: the stages an agent moves through from draft to retirement
MCP	Model Context Protocol: a standard way for agents to connect to tools and data
RAG	Retrieval-augmented generation: answering from your own documents, with citations
Computer use (CUA)	An agent that operates software through its visual interface, like a person would
Feature flag	A switch that turns an agent or change on for some users before everyone

SOURCES

IBM Introduces Identity and Control for AI Agents, IBM Newsroom, October 27, 2026

Gartner Identifies Six Steps to Manage AI Agent Sprawl, Gartner, April 28, 2026

StackAI. Secure, governed AI agents, ready for production.

STACKAI.COM

BUILD AGENTS WITH GOVERNANCE BUILT IN

Book a demo

In a 30-minute walkthrough, see how your teams can build agents on StackAI with governance applied from the first step.

Book a Demo →

Built for your workflow

See an agent for a process like yours, such as an IT support desk or KYC review.

Governed from day one

Owners, permissions, approval gates and run logs, shown live in the product.

Deployed on your terms

Our cloud, your VPC, on-premise or air-gapped, with SSO and role-based access.