



Agentic AI Governance Fundamentals

How to go from shadow AI to verifiable
trust for agentic deployments

eBook | 2026



Introduction:

The AI Governance Challenge

Agentic AI sits poised to transform the way organizations do business, promising efficiency gains and new, innovative ways to empower employees and engage customers. But the security risks are as significant as the promised rewards.

AI agents can reason, plan, and execute actions autonomously. They operate across environments and trust domains, and they create dangerous security gaps if given too much leeway.

In short, AI agents need governance.

Nearly 80% of organizations report that AI adoption is already outpacing their current governance capabilities, and only 11% believe they are fully ready for the scale of AI agent deployment planned for the next year.¹

The good news is that the tools needed to govern AI already exist. In this eBook, you'll see how DigiCert® AI Trust Manager can help organizations build a solid foundation of trust for agents, while simplifying AI governance and enabling AI scaling with confidence.




Agent discovery and inventory

Before agents can be governed, organizations need to know they exist. This isn't as simple as it sounds. Today, 98% of organizations have employees using unsanctioned AI.² Shadow AI is a universal problem, and too often it is the IT and security teams themselves that build and deploy new agents. The first step to governing agents must be discovery. The discovery process should produce a clear inventory of all existing agents so that they may be effectively governed. Because new agents are being created every day, the discovery process should repeat regularly to keep the inventory list accurate.

AI Trust Manager allows organizations to automatically discover agents running in defined agent environments such as AWS or GCP, building a centralized private registry. Organizations can also choose to publish trusted AI assets to the DigiCert public registry, which is built on the Agent Name Service (ANS) Registry standard.

Having a complete registry of current agents allows administrators to proceed with identity, policy, and lifecycle management.



Agent identity

Once organizations have visibility into the agents running across their environments, these agents need identity.


Traditional IAM defaults to long-lived credentials and relies on MFA, passwords or biometrics. It is not suitable for agents, particularly when they might be spun up and down quickly and frequently. Long-lived and static credentials can let a compromised or manipulated agent access systems indefinitely at machine speed, potentially turning a single compromised token into a multi-system security breach. Even agents that have not been manipulated by bad actors might misuse stored secrets to access and use sensitive data in unexpected ways.

To eliminate these risks, what agents really need are runtime attestation and ephemeral credentials that are cryptographically verifiable and continuously validated.

For the agents built in-house by the organization, AI Trust Manager issues short-lived SPIFFE identities (SVIDs), using short-lived credentials that are automatically rotated, eliminating the risks inherent in shared API keys and long-lived static credentials. This enables mTLS authentication between agents and systems.

For third party agents, such as Microsoft Copilot or ServiceNow, DigiCert embeds a Model Context Protocol (MCP) gateway within the agent's toolchain, which evaluates the agent's identity and policy when the agent takes an action.


In both cases, identity must be tied to policy.



Policy and the Agent Passport


When it comes to policy and authorization, organizations need a clear way to control what systems and data agents can access, as well as what actions they are allowed to perform. All this information, alongside the agent's identity, can be encoded into a single governance artefact.

This is the DigiCert AI Passport: a cryptographically signed capability and governance document bound directly to an agent's workload identity for built agents, and agentic sessions for SaaS agents.



The Passport can encode a variety of important information about the agent, including:

- The agent's bound SPIFFE workload identity
- The systems and MCP connections it is authorized to access
- The operations it is permitted to perform
- The environments in which it may operate
- The sensitivity level of data it is approved to handle
- The policies that govern its behavior
- The expiration and renewal lifecycle of its authorization
- The organization that owns the agent
- The named accountable human owner responsible for it




Just as a real-world passport is issued by a sovereign authority, the AI Agent Passport is issued and signed within a defined trust domain. Its authority is rooted in cryptographic validation, and it can be used across clouds, platforms, and organizational boundaries.

Agent lifecycle management


Much like real-world passports, AI Agent Passports have a lifecycle.



Passports can expire.




Passports can be renewed



Most critically, passports can be revoked.

Recent news stories about rogue agents behaving badly have inspired significant concern and even “kill switch” legislation.³ Autonomous agents operate at machine speed and scale, and trust conditions can change due to any number of factors such as policy violations, risk posture degradation, or administrative action. When these changes occur, privileges should be easy to revoke.

AI Trust Manager supports workflows for Passport issuance, renewal, expiration, and revocation, covering the entire agent lifecycle. To simplify lifecycle management, these workflows can be automated. Passports can be issued and renewed by the system based on organizational settings. The kill switch is also automated, triggering based on agent actions to revoke the passport.



Agentic session trust

With discovery, inventory, identity, authorization, and lifecycle management in place, the foundation for governing AI agents is nearly complete.

To bring it all together, organizations need a way to track and protect agent sessions. AI Trust Manager provides a continuously updated relationship graph linking agents to systems, models, connections, and data classifications. This enables proactive blast-radius analysis, revealing the impact of revoking a Passport or changing a policy, and allows for improved vulnerability detection. Crucially, associating lineage with data sensitivity also allows for the enforcement of additional controls such as data masking, access segregation, or environment-based restrictions. Compliance frameworks like HIPAA or GDPR become enforceable through technical controls rather than documentation alone.

Without lineage, authorization operates in isolation. With lineage, governance becomes systemic rather than reactive.

The screenshot displays the 'demo-mcp-server' agent page in the AI Trust Manager. At the top, the agent is labeled 'demo-mcp-server' with status indicators 'Deprovisioned' and 'Not Running'. Below this, it shows ownership and update information: 'Owned by admin@digicert.local · Updated by admin@digicert.local · Aug 19, 2026, 03:52 PM'. The agent's metadata includes 'NAMESPACE: wl-fpucsb', 'CLUSTER: dev-gcp-test', and 'ID: fpucsbekhtbv19n01v3wif'. It also shows the 'ENFORCED VERSION' as 'v0.2.13 → v0.2.17' with an 'Upgrade' link.

A yellow warning box indicates '2 ACTIVE CONDITIONS':

- enforcement_not_reporting**: the PEP sidecar has stopped reporting enforcement counters (5 days ago)
- workload_destroyed**: workload is destroyed

The 'Activity' tab is selected, showing a table of recent activity:

RECENT ACTIVITY		All Violations	All Logs
Last 5 policy denials for this agent.			
Constraint Deny	rate_limited api.openai.com		2 days ago
Constraint Deny	rate_limited demo-mcp-server.wl-fpucsb.svc.cluster.local/mcp/v1/tools/*/call		2 days ago
Invalid Identity	svid_not_available hooks.slack.com/services/*		2 days ago
Invalid Passport	invalid_spiffe_mismatch api.openai.com/v1/chat/completions		2 days ago

Result: AI governance that scales

AI agents operate across boundaries at machine speed. They require:



Clear visibility



Verified identity



Scoped authorization




Named ownership



Enforceable validation points



Lifecycle management



Traceable movement

Many emerging solutions in AI governance focus on prompt inspection, anomaly detection, or post-execution monitoring. While these capabilities are valuable, they do not establish foundational trust. The industry will ultimately require structural controls: verifiable identity, signed capability authorization, enforceable policy evaluation, lifecycle management, lineage-aware impact analysis, and trust that spans agentic sessions.

By applying identity controls to agents, binding them to signed Passports, enforcing policy through defined control points, and mapping their lineage across systems and data types, enterprises can scale autonomy without sacrificing control.

At DigiCert, we believe that the future of enterprise AI will not be defined solely by how intelligent systems become, but by how deliberately we govern them. This is why we've built an AI Trust solution on the proven principles of PKI and DNS, as well as the AI Root of Trust that forms a solid and reliable foundation for organizations seeking to gain control of agent sprawl and scale agentic AI with confidence.

Ready to see DigiCert AI Trust Manager in action?



[Try the interactive demo today.](#)

About DigiCert

DigiCert is a global leader in intelligent trust. We protect the digital world by ensuring the security, privacy, and authenticity of every interaction. Our AI-powered DigiCert ONE platform unifies PKI, DNS, and certificate lifecycle management, to secure infrastructure, software, devices, messages, AI content and agents. Learn why more than 100,000 organizations, including 90% of the Fortune 500, choose DigiCert to stop today's threats and prepare for a quantum-safe future at www.digicert.com.

© 2026 DigiCert, Inc. All rights reserved. DigiCert is a registered trademark of DigiCert, Inc. in the USA and elsewhere. All other trademarks and registered trademarks are the property of their respective owners.

The agent governance lifecycle

1 | Discover



Identify every AI agent across your enterprise and environments

2 | Register



Register agents and capture metadata for identity, policy, and control

3 | Identify



Issue AI Agent Passports that bind agent identity, metadata, policies, and capabilities

4 | Govern



Define policies that govern agent behavior and control access to connected MCP tools

5 | Enforce



Enforce policy with DNS and/or PKI for validation and security

Citations

1. <https://newsroom.ibm.com/2026-06-08-new-ibm-study-finds-cios-and-ctos-face-growing-ai-control-gap-as-enterprise-deployment-scales>
2. <https://www.netskope.com/netskope-threat-labs>
3. <https://www.cnn.com/2026/08/06/ai-kill-switch-bill-openai-anthropic-meta.html>