Where Okta for AI Agents helps meet NIST's design tenets for agent identity

How Okta for AI Agents maps to the design tenets in NIST's post, Design Tenets for Agent Identity Management Systems

Executive summary

An AI agent can be steered by what it reads. An instruction hidden in a document, web page, or email can send an agent with broad access to forward files outside the company or call tools it was never meant to use. NIST's new post, Design Tenets for Agent Identity Management Systems1, says the fix cannot live inside the model: "Guardrails alone are insufficient."

Large Language Models (LLMs) receive the instructions they are given and the data they work with as a single stream of text, so they cannot reliably distinguish between the two. NIST concludes that this separation has to be built into the architecture. The system that determines what an agent is allowed to do must sit outside the agent, between the agent and the data, tools, and credentials it uses.

NIST published the post on October 27, 2026. Its ten design tenets come from public comments on the NCCoE's project on software and AI agent identity and authorization2, and NIST asks organizations to apply them in proportion to the risk of the transactions they support. The authors call them an initial set. They say some of the tenets may not be fulfilled by the current state of standards and technology, and they offer them as a way to assess current capabilities and build a roadmap.

Okta for AI Agents at a glance

Okta for AI Agents gives AI agents their own identity with a named owner and issues short-lived tokens scoped to the task. It enforces policy at Agent Gateway, outside the agent, decides access at request time, reviews that access on a recurring schedule, logs tool calls, and lets security teams stop an agent with a single action. The sections that follow map these capabilities to NIST's tenets. NIST notes that not every tenet can be met with today's standards and technology. Against that bar, Okta for AI Agents covers eight tenets today, partly covers tenet 6 (signed mandates), and has tenet 5 on the roadmap. The sections show what the product does against each tenet, not a claim to meet every clause of each one.

1. https://www.nist.gov/blogs/cybersecurity-insights/design-tenets-agent-identity-management-systems

2. https://pages.nist.gov/nccoe-ai-identity/

© 2026 Okta, Inc. and/or its affiliates.

Tenet 1: Distinct agent identification

NIST asks for agents with an identity, credential, and access management lifecycle distinct from human principals. A unique identity helps prevent an agent from impersonating a human user or service account and traces autonomous actions back to a specific software component.

Okta for AI Agents addresses this in three ways:

Register homegrown agents: The AI agent registry adds custom-built agents to the Universal Directory as first-class identities, with a named human owner, lifecycle policies, and their own credential (private key JWT, registered directly or through a client ID metadata document), separate from people and service accounts.

Import agents from external sources: Agent import brings in agents from Salesforce Agentforce, Amazon Bedrock AgentCore, ServiceNow AI Platform, and other platforms, so they fall under the same policy as the ones you built.

Find shadow agents: The Okta browser plugin detects shadow AI agents by identifying OAuth consent grants in managed Chrome browsers. EDR integrations surface coding agents, such as Claude Code and Cursor, on managed devices and tie them to named owners.

Tenet 2: Agent linking, delegation, and attenuation

NIST asks that an agent and whoever initiated its task stay linked throughout execution, and that authority be narrowed as it passes to sub-agents or downstream services, so that no downstream party can claim more than its parent. Without that link, high-impact actions cannot be attributed to the responsible party.

Okta for AI Agents addresses this in three ways:

Name an owner: Agents registered in Universal Directory can have a named human owner.

Mediate agent-to-app access: Cross App Access, built on the open IETF ID-JAG standard, lets your identity provider mediate agent-to-app authorization, so employees are not asked for a fresh consent per agent.

Control agent-to-agent handoffs: Agent-to-agent connections let agents call other agents under an admin-approved policy, with handoffs recorded. This helps limit an agent's ability to gain access by handing work to one with more.

© 2026 Okta, Inc. and/or its affiliates.

Tenet 3: Least and zero standing privilege

NIST asks that agents hold no persistent rights when idle, that access be bounded to the minimum scope a task requires, and that agents be kept from combining privileges that conflict (separation of duties). An agent with standing or broad privileges can take unintended actions or be hijacked to reach sensitive data.

Okta for AI Agents addresses this in three ways:

Scoped, short-lived tokens: Agents get tokens scoped to what a task needs, for as long as it needs it, across authorization servers, vaulted secrets, service accounts, applications, and MCP servers. This helps reduce standing access and credentials embedded in code. Vaulted secrets and service accounts are not part of Okta for AI Agents - Core.

Credentials the agent does not hold: Agent Gateway injects the credential at call time, so the agent does not need to hold it.

Recurring access reviews: Access certifications enable agent owners to review and revoke both who can reach an agent and what the agent can reach, on a recurring cadence, with an audit trail.

Tenet 4: Short-lived, dynamic, and revocable authorization

NIST asks for grants that are just-in-time, time-bound, and audience-specific, with tool calls and resource access verified at runtime and authority revocable when anomalous behavior or context drift appears. NIST says to avoid long-lived bearer tokens and static, reusable secrets, because a stolen credential becomes persistent access and a valid token cannot be interrupted if an agent is hijacked mid-task.

Okta for AI Agents addresses this in two ways:

Short-lived tokens at the gateway: At Agent Gateway, the agent receives a short-lived token scoped to the gateway, and Okta checks tool calls.

Disable and revoke agent access: One action in Universal Directory deactivates an agent, revokes its access to linked resources within the Okta security boundary, and blocks new token requests.

© 2026 Okta, Inc. and/or its affiliates.

Tenet 5: Human-in-the-loop triggers

NIST asks organizations to set risk thresholds that trigger human review, to send the right person enough context to decide, and to watch for consent fatigue, where escalation becomes so frequent that people approve by habit. Without enforced triggers, a malfunctioning or manipulated agent can do irreversible damage.

Okta for AI Agents addresses this in two ways:

Risk-based approval (Roadmap): Human-in-the-Loop for AI Agents is designed to add approval by role, action type, or risk level.

Fewer routine prompts: Cross App Access helps reduce the consent prompts employees see for routine agent access, leaving approvals for actions that require judgment.

Tenet 6: Mandates and constraints

NIST asks that task originators issue signed mandates that state intended outcomes and limits, which downstream services can verify.

Okta for AI Agents addresses this in part:

Scoped tokens: Agents get tokens scoped to what a task needs, which bounds what an agent can do for that task.

Cross App Access: Your identity provider mediates agent-to-app authorization, so policy sets the scope an agent receives, not the agent.

Signed task-level mandates (Roadmap): Task-level intent, expressed through standards such as RAR and transaction tokens, is on the roadmap.

© 2026 Okta, Inc. and/or its affiliates.

Tenet 7: Architectural separation of policy from reasoning

"Guardrails alone are insufficient." NIST, Design Tenets for Agent Identity Management Systems

NIST asks that the logic deciding what an agent attempts be architecturally separate from the systems deciding what it is allowed to do, so that a prompt injection or malicious data processed by the model cannot manipulate credentials or expand the agent's permissions.

Okta for AI Agents addresses this in two ways:

Policy at the gateway: Agent Gateway is a single endpoint that agents call over MCP to access enterprise tools, regardless of which vendor built the agent. Okta verifies the agent, checks which tools it may use, injects the credential, and logs the call against a managed identity. The policy lives in the gateway, outside the agent, helping keep the text the model reads from widening its permissions. The same separation holds when the enforcement point is a partner gateway such as Amazon Bedrock AgentCore, where Okta still issues the scoped token the gateway enforces.

Authorization outside the agent: Fine-grained authorization keeps authorization logic out of the agent. Policy changes take effect at the next request with no change to agent code.

Tenet 8: Deterministic access decisions

NIST asks organizations to keep deterministic mechanisms as the primary basis for access decisions. Probabilistic information, such as a risk score, can feed the decision but does not make it. That limits the effect of prompt injection and makes agent behavior more predictable when it reaches for sensitive services, data, and tools.

Okta for AI Agents addresses this in two ways:

Scopes set the boundary: OAuth scopes and Cross App Access fix what an agent can reach before it makes any request, so the decision rests on granted access, not on anything the model reads or a risk score.

Decisions at request time: Fine-grained authorization then evaluates current relationships and conditions at the moment of the request, not roles assigned months ago.

© 2026 Okta, Inc. and/or its affiliates.

Tenet 9: Continuous monitoring

NIST points out that agents operate at a scale and velocity beyond humans and most deterministic automation, so systems need agent-specific metrics, thresholds, and policies to detect anomalous behavior. A missed change in an agent's risk posture can leave compromise, configuration drift, or privilege escalation undetected.

Okta for AI Agents addresses this in two ways:

Find the agents: Shadow agent discovery surfaces agents that were not registered, through the browser plugin and EDR integrations.

Record what they do: Audit logs and telemetry capture tool calls, access attempts, and authorization decisions, and stream to your SIEM.

Tenet 10: Immutable evidence

NIST asks for a tamper-evident, append-only record of agent identity, delegated intent, and actions, stored independently of the agent.

Okta for AI Agents addresses this in three ways:

Immutable audit trail: System Log records agent token exchanges, configuration changes, and authorization decisions in an immutable audit trail.

Outside the agent's control: The record is kept by Okta, not the agent, and streams to your SIEM for retention.

Recurring access audits: Access certifications put agent access in front of its owners on a recurring schedule. Each review and revocation decision is recorded, so you have the runtime log and an ongoing audit of who can reach what.

© 2026 Okta, Inc. and/or its affiliates.

Built on open standards

Cross App Access is built on the IETF Identity Assertion JWT Authorization Grant, an open, vendor-neutral protocol that Okta authored. Any agent or app that supports the standard can take part, not only Okta's. Agent SSO, included in the workforce SSO plan, applies it to any agent that supports it.

One attack, four tenets

NIST notes that multiple tenets may need to be in place to mitigate a specific attack. Its example is the confused deputy, where a prompt injection convinces an agent with broadly scoped access to execute an action it otherwise would not. Stopping it requires linking and delegation (tenet 2), least privilege (tenet 3), separation of policy from reasoning (tenet 7), and deterministic access decisions (tenet 8) to work together. NIST's point is that one tenet on its own does not close it.

Three questions to ask

Use the tenets as a review checklist the next time an agent platform comes up in a security review.

1. Can anything the agent reads change what it is allowed to do?
2. Does the agent ever hold the credential for the system it calls?
3. How fast can you stop one agent, and what stops with it?

Vendors that answer the first one with a better prompt or a smarter filter are describing a guardrail. NIST's position is that the answer has to be architectural.

How Okta can help

Okta for AI Agents gives AI agents a managed identity and enforces policy on what they can access, whether teams build the agents or adopt them, and regardless of which identity provider you run for your workforce.
It organizes the work around four questions:

Where are my agents? Tenet 1: a unique identity with a named owner for agents, and discovery of agents nobody registered.
What can they do? Tenets 2, 3, 6 (in part), 7, and 8: scoped, short-lived access, policy enforced outside the agent, and decisions made at request time.
What are they doing? Tenets 7, 9, and 10: a logged, immutable record of tool calls, streamed to your SIEM.
How do I respond? Tenets 4 and 5: one-action deactivation and risk-based approval (Roadmap).

For third-party agents such as Claude or Cursor, Agent SSO is included in workforce SSO and gets them into the directory with their own identity. For the agents you build, tenets 3, 4, 7, 8 and 9 above are Okta for AI Agents.

Learn more at okta.com/products/govern-ai-agent-identity.

About Okta

Okta is The World's Identity Company™. We secure Identity, so everyone is free to safely use any technology. Our customer and workforce solutions empower businesses and developers to use the power of Identity to drive security, efficiencies, and success — all while protecting their users, employees, and partners. Learn why the world's leading brands trust Okta for authentication, authorization, and more at okta.com.

Any mention in this white paper or datasheet of solutions, features, functionalities, certifications, authorizations, or attestations that are not currently generally available or have not yet been obtained may not be delivered or obtained on time or at all. We assume no obligation to deliver on such items and you should not rely on them to make your purchase decisions.

These materials are for general informational purposes only and do not constitute legal, privacy, security, compliance, or business advice. The content may not reflect the most current security, legal and/or privacy developments. You are solely responsible for obtaining advice from your own legal and/or professional advisor and should not rely on these materials. Okta makes no representations or warranties regarding this content and is not liable for any loss or damages resulting from your implementation of these recommendations. Information on Okta's contractual assurances to its customers may be found at okta.com/agreements.

NIST has not reviewed or endorsed Okta products.

© 2026 Okta, Inc. and/or its affiliates.