IDC PERSPECTIVE

The Domain Is the Root of Trust: DNS-Based Agent Authentication

Frank Dickson Christopher Rodriguez Grace Trinidad

EXECUTIVE SNAPSHOT

AI agents are proliferating faster than the governance frameworks meant to contain them. Most enterprises cannot answer basic questions about which agents are operating in their environments, what these agents are authorized to access, or how to stop one that goes rogue. DigiCert is positioning DNS, the protocol that already underpins internet-scale trust, as the enforcement layer for agent authentication and access control. This IDC Perspective examines the technical approach, its parallels to proven email authentication standards, and what enterprise security teams should consider before deploying it.

Key takeaways

- AI agents are proliferating rapidly, outpacing current governance and security frameworks, leaving enterprises unable to track, authorize, or control agent activity effectively.
- DigiCert proposes leveraging DNS as the enforcement layer for agent authentication and access control, drawing on proven models like DMARC for email, to provide scalable, domain-based trust without requiring new endpoint software.
- The DNS-based approach enables pre-connection blocking of unauthorized agent activity, with policy records specifying permitted actions and a kill switch that terminates sessions if policy checks fail, preventing exploits before they start.
- While the model is innovative and low disruption, enterprises must ensure DNS policy records remain current and accurately reflect agent inventories and must rigorously test policy enforcement and kill switch mechanisms to avoid misconfiguration and security gaps.

Recommended actions

- Evaluate DNS-based agent access controls as a near-term, low-disruption strategy for agentic AI governance, prioritizing this approach over more disruptive infrastructure changes that require hardware or platform migration.

June 2026, IDC #US54567926

- Establish and maintain a comprehensive inventory of all AI agents — both sanctioned and unsanctioned — operating within the environment prior to deploying DNS policy records for agent scope enforcement, ensuring policy coverage is complete and current.
- Align agent identity management with workload identity frameworks (e.g., IETF WIMSE, NIST CSF 2.0), treating agents as governed workloads requiring runtime attestation and short-lived credentials, rather than extending traditional human IAM systems.
- Require vendors to demonstrate real-time OPA policy engine updates and validate the effectiveness of the DNS-based kill switch mechanism through simulated prompt-injection and domain-spoofing scenarios before production deployment, ensuring robust enforcement and incident response capabilities.

SITUATION OVERVIEW

AI agents are not well-behaved software. AI agents are designed to be highly autonomous but with no accountability. The combination has led to unpredictable, often disastrous results for some organizations that adopt AI agents without implementing proper guardrails. Unlike conventional applications that execute a fixed set of instructions against known endpoints, agents reason, improvise, and spawn subagents and traverse trust boundaries without pausing to ask permission. Although most organizations are deploying AI agents, IDC has identified inadequate governance and security controls as the leading reason agentic AI projects fail.

Traditional IAM controls do not transfer cleanly to agents. Traditional platforms were designed for users who respond to MFA prompts and authorize consent screens. Agents cannot. The fallback is static API keys, which represent exactly the kind of long-lived, broadly scoped credentials that security teams are actively trying to eliminate from their environments.

As a result, there is a shift toward treating agent identity as a workload identity problem rather than an extension of human IAM. Emerging approaches align with frameworks such as IETF WIMSE and NIST CSF 2.0, in which agents are issued short-lived credentials and require runtime attestation.

DMARC was here first: The DNS precedent

DigiCert asserts that the answer to the agent access problem lies in a mechanism enterprises already trust and operate at scale: DNS. The insight is structural. Every agent action, whether downloading a package, connecting to an MCP server, or calling an API, begins with a DNS query. DNS is therefore not just a naming service; it is the natural chokepoint for agent access control.

The model DigiCert proposes maps directly onto DMARC, the email authentication framework that solved domain impersonation by tying sender identity to domain ownership. In the DMARC model, a sending organization:

- Publishes authorized mail servers via SPF records in DNS
- Cryptographically signs outbound email with keys referenced in DNS via DKIM
- Publishes a policy record instructing receivers to reject messages that fail those checks

A receiver querying **_dmarc.ibm.com** can verify whether a message actually originated from IBM before it reaches an inbox.

The agent analog works the same way, with the domain as the root of trust. An organization publishes authorized agent identities and their issuing certificate authority (CA) in a DNS TXT record. This record also declares permitted scopes, read-only access, specific APIs, and an allowlisted set of domains. A policy record tells receiving gateways what to do when an agent's credential or scope declaration fails: terminate the session. A gateway querying **_agentpolicy.ibm.com** can verify whether an inbound agent claiming to originate from IBM is legitimate before it is admitted to the environment. The approach requires no new endpoint software, leverages DNS infrastructure that already exists, and follows the same adoption path DMARC took from concept to RFC 7489.

DigiCert implements this through its platform, delivered on top of its UltraDNS infrastructure. The technical flow has four stages. An external SaaS agent arrives at an MCP gateway presenting a credential that claims a domain of origin. The gateway queries DNS for the agent's claimed domain. The gateway then validates the credential against three conditions: whether it was signed by the DigiCert CA, whether the requested scope falls within what the DNS record permits, and whether the agent's JWT passport is unexpired and valid. This passport artifact is the DigiCert answer to agent policy and authorization. It is based on open standards and validated at runtime for authenticity, expiration, and scope before access is granted. If all three pass, the agent operates through the MCP gateway, with approved tools visible and OPA policies that scope access to declared resources only. If any check fails, the DNS query is blocked, the MCP gateway kills the session, and the event is logged.

The kill switch is not a secondary safeguard. It is central to the design. If an agent contacts an unauthorized domain at any point during execution, DNS blocks the query before the TCP connection is established. The MCP gateway then terminates the agent session and fires an alert. This pre-connection blocking is architecturally significant: The malicious domain never receives a request, breaking the exploit chain that begins with a successful connection before it starts.

As part of its broader platform capabilities, DigiCert also provides mechanisms that can support foundational prerequisites for this model, including discovery and inventory of machine identities and certificates across environments, the issuance of cryptographically verifiable workload identities to agents, and the issuance of agent passports. These capabilities can be extended to improve visibility into AI agents and their associated credentials.

ADVICE FOR THE TECHNOLOGY BUYER

The DNS enforcement model is genuinely novel and grounded in proven technology. DNS-based policy verification has operated at internet scale for over a decade in email. Porting it to agent authentication is conceptually coherent, and the zero-endpoint software deployment model addresses a real objection security teams face when introducing new controls into complex environments.

That said, enterprises should evaluate several open questions before committing. First, the model assumes that DNS policy records remain accurate and current as agent populations evolve. Organizations managing dozens of sanctioned agents today may be managing hundreds within 18 months. Keeping DNS TXT records in sync with a rapidly changing agent inventory requires operational discipline that most security teams have not yet developed. Stale records create gaps that are difficult to detect until they become incidents.

Moreover, while the current assumption is that AI agents willingly declare their status and connect via proper MCP channels, this may change over time. The broad adoption of DNS-based policy enforcement may be seen as undesirable friction by agent operators. This may lead to efforts to sidestep controls through evasive or deceptive means, such as impersonating a human or other, more popular agent. Cybersecurity trends are often cyclical, with adaptations on both sides of the equation. DNS-based policy enforcement, while an important foundational starting point, may eventually require complementing support by other web traffic monitoring and profiling technologies.

Second, the scope declaration mechanism, in which a DNS record specifies permitted actions such as read-only access or access to specific APIs, introduces a powerful but fragile policy-as-infrastructure pattern. An overly permissive scope record functionally defeats the control. The OPA policy engine that governs enforcement requires tuning and maintenance; it is not a set-and-forget component.

Essential guidance

- Evaluate DNS-based agent access controls as a near-term, low-disruption first step toward agentic AI governance, prioritizing it ahead of deeper infrastructure changes that require hardware or platform migration.

- Before deploying DNS policy records for agent scope enforcement, build an inventory of all agents currently operating in your environment — both sanctioned and unsanctioned — since a policy model is only as good as the completeness of what it covers.
- Treat agent identity as a workload identity problem rather than extending existing human IAM systems, and align procurement and architecture decisions to IETF WIMSE and NIST CSF 2.0 guidance, which frame agents as governed workloads requiring runtime attestation and short-lived credentials.
- Require vendors implementing DNS-based agent trust to demonstrate how their OPA policy engine handles scope updates in real time, and test the kill switch mechanism under simulated prompt-injection and domain-spoofing scenarios before production deployment.
- Audit the gap between agents you build and agents you procure, since SaaS agents from vendors such as Salesforce, Microsoft, and Anthropic require a distinct identity and policy pathway that must be kept consistent with internally built agent governance to avoid creating blind spots.

LEARN MORE

Related Research

- *RSAC Conference 2026: Trust But Verify — Continuously, Identity at the Forefront of Nondeterministic Security* (IDC #IcUS54481626, April 2026)
- *RSAC Conference — Human-in-the-Loop* (IDC #IcUS54470726, April 2026)
- *Microsoft Agent 365: The Control Plane for Operationalizing Agents for Day 2 and Beyond* (IDC #EUR154499226, April 2026)
- *Identity for AI Agents: From Gatekeeper to Command Center* (IDC #US54357626, March 2026)
- *From Technology Provider to Orchestrator of Autonomy: The CIO's Framework for Industrializing Agent Vetting* (IDC #US54307326, February 2026)

Synopsis

This IDC Perspective explores DigiCert's approach to AI agent authentication using DNS as the enforcement layer, drawing parallels to established email authentication standards like DMARC. By leveraging DNS infrastructure, organizations can control agent access and enforce policies without new endpoint software. The model promises scalable, preemptive security but requires disciplined management of DNS records and careful policy configuration to avoid gaps. Enterprises are advised to treat agent identity as a workload issue and prioritize DNS-based controls for near-term AI governance.

"The agent inventory problem is the new shadow IT problem, except the blast radius is larger and the audit trail is thinner. DNS-based policy enforcement is one of the few controls that can scale as fast as the agents themselves," said Frank Dickson, group vice president, Security and Trust at IDC.

ABOUT IDC

International Data Corporation (IDC) is the premier global market intelligence, data, and events provider for the information technology, telecommunications, and consumer technology markets. With more than 1,300 analysts worldwide, IDC offers global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries. IDC's analysis and insight help IT professionals, business executives, and the investment community make fact-based technology decisions and achieve their key business objectives.

Global headquarters

One Beacon Street
Suite 33100
Boston, MA 02108
USA
508.872.8200
X: @IDC
blogs.idc.com
www.idc.com