The Agentic Identity Stack: Where Okta, Microsoft Entra, Astrix, Oasis, and ScrambleID Fit Together

Buyer's Guide / Last updated 2026-06-10 / https://www.scrambleid.com/learn/agentic-identity-stack

Buyers keep asking this as a versus question. It's an architecture question.

TL;DR (canonical)

- Agentic identity isn't one product category. It's a **stack with three layers**: agent directories and lifecycle governance (Okta for AI Agents, Microsoft Entra Agent ID), agent and NHI discovery, posture, and detection (Astrix Security, Oasis Security), and **per-action proof**: the cryptographic record that a specific agent, under a specific authority, performed a specific action (ScrambleID).
- The layers **compose rather than compete**. Discovery tools feed the directory; the directory governs the lifecycle; the proof layer signs and evidences each action the governed agents take.
- ScrambleID doesn't replace your IdP or your posture tooling. It adds the cryptographic spine they don't claim to deliver: a signature in every request, per-action authority, and a customer-verifiable, non-repudiable audit trail.
- With Gartner expecting **40% of enterprise applications to ship task-specific AI agents by the end of 2026** (up from under 5% in 2025), most enterprises will run more than one layer of this stack. The design question is what each layer must answer, not which logo wins.

Why one product doesn't cover it

An agent estate generates three different obligations, usually owned by three different teams. Someone has to **know what exists**: every agent, MCP server, and non-human identity, sanctioned or not. Someone has to **govern the lifecycle**: registration, ownership, scoping, deprovisioning, conditional access. And someone has to **prove what happened**: when an agent files, pays, approves, or deletes, the record has to hold up in front of an auditor, a regulator, or a customer dispute.

These obligations show up directly in the frameworks. The OWASP Top 10 for Agentic Applications names identity and privilege abuse (ASI03) and rogue agents (ASI10); a February 2026 NIST NCCoE concept paper calls for every agent action to be tied back to an accountable identity, auditable and

non-repudiable. No single layer of the stack answers all of that alone, which is why the market grew three layers in the first place.

Layer 1: agent directories and lifecycle governance

Okta for AI Agents treats agents as first-class identities in Universal Directory: discovery of known and shadow agents (including through OAuth consent-grant detection), onboarding with a named human owner, short-lived credentials in place of long-lived tokens, lifecycle workflows, a kill switch, and an audit trail. Okta is also championing open standards for the category, ID-JAG and Cross App Access, and has extended the capability set beyond its own IdP.

Microsoft Entra Agent ID extends Entra to agent identities: purpose-built agent identities created from blueprints with parent-child relationships, OAuth, MCP, and A2A support, Conditional Access and risky-agent detection, governance workflows that keep agents from being orphaned, and agent sign-in and audit logs. It reaches third-party runtimes through workload identity federation and an SDK sidecar.

If your enterprise runs on either ecosystem, this layer is where agent registration and lifecycle live. That's the right place for it: the directory is where your humans, groups, and policies already are.

Where ScrambleID composes with this layer: the directory governs who the agent is and whether it may operate. ScrambleID adds proof of what it did: each call carries the agent's signature, generated at the point of action, so the directory's audit trail gains a cryptographic record beneath it. We don't replace your IdP. We add the spine of evidence under the lifecycle it manages.

Layer 2: discovery, posture, and detection

Astrix Security inventories AI agents, MCP servers, and NHIs in real time, including shadow and unregistered agents, remediates excessive privileges and risky configurations, detects anomalous agent behavior, and provisions agents secure-by-design through its Agent Control Plane with just-in-time, precisely scoped credentials.

Oasis Security runs the NHI Security Cloud and Agentic Access Management: real-time inventory with ownership mapping, posture management across cloud and SaaS, lifecycle automation, anomaly detection, and access control built around agent intent rather than static roles.

This layer answers the question the shadow-agent problem makes urgent: what's actually running, who owns it, and is its posture sane. Enterprises with sprawling multi-cloud NHI estates tend to feel this pain first, and these platforms are built for it.

Where ScrambleID composes with this layer: discovery tells you what exists; posture tells you what's risky. ScrambleID gives the agents that survive that triage their operating primitive: an identity with zero static secrets and a signature on every action. The fewer long-lived credentials exist, the

smaller the attack surface the posture layer has to watch, and the inventory stops churning against a sea of unattributable keys.

Layer 3: per-action proof

This is the layer ScrambleID builds, and it starts where the other two stop: at the individual action. ScrambleID's non-human identity gives every agent its own cryptographic identity with zero static secrets, and Per-Action Authority adds the control surface for the actions that matter: a signature on every action, a human cosigner wherever policy demands one, and a hash-chain ledger the customer can verify independently.

Three properties define the layer. The verification is **independent of the platform the agent runs on**, so the evidence doesn't depend on the runtime it describes. The keys are **device-bound or customer-held**, so the proof belongs to you. And the result is **non-repudiable**: when the log says an agent did something, the signature in the request is the evidence, not the log line's word for it.

For an auditor, the three layers answer three different questions. The inventory answers "what exists." The directory answers "what was it allowed to do." The proof layer answers "what did it actually do, and can you prove it." Run together, the answers chain.

The stack at a glance

Layer	What it answers	Platforms there
Agent directory and lifecycle governance	Who is this agent, who owns it, what may it do, when does it expire?	Okta for AI Agents, Microsoft Entra Agent ID
Discovery, posture, detection	What exists (including what nobody registered), what's risky, what's behaving abnormally?	Astrix Security, Oasis Security
Per-action proof and authority	What did it actually do, under whose authority, and can anyone verify that independently?	ScrambleID

Key Takeaway

The agentic identity conversation reads like a bake-off and behaves like a stack. Directories govern, discovery platforms find and watch, and the proof layer evidences. ScrambleID is built for the third layer and built to compose with the first two: your IdP keeps the lifecycle, your posture tooling keeps the watch, and every action your agents take picks up a signature that holds.

FAQ

Does ScrambleID replace Okta for AI Agents or Microsoft Entra Agent ID?

No. The directory layer is where agent registration, ownership, and lifecycle belong, and ScrambleID deploys as an overlay alongside your existing IdP. What it adds is the per-action proof those platforms don't claim: a cryptographic signature in each request and a customer-verifiable ledger of what every agent actually did.

We already run Astrix or Oasis. What does ScrambleID add?

Discovery and posture tooling answers what exists and what's risky. ScrambleID changes what the discovered agents run on: per-agent identity with zero static secrets and a signature on every action. The two are direct complements; a clean inventory full of long-lived keys still can't prove which agent performed which action.

Which layer should we buy first?

Sequence by your sharpest pain. Unknown estate: start at discovery. Governance findings or framework pressure: start at the directory. Agents executing consequential actions (payments, filings, approvals, deletions): start at proof, because that's where an incident becomes a liability question. Most enterprises end up with more than one layer; the architecture above is how they avoid overlap.

How does this stack map to the agentic-governance frameworks?

The OWASP Agentic Top 10's identity items (ASI03, ASI10) span all three layers: discovery finds the rogue agents, the directory scopes privileges, and per-action signatures make abuse attributable. The NIST NCCoE concept paper's accountability language (every action tied to an identity, auditably and non-repudiably) is the proof layer's job description. The full mapping lives in our compliance article.

Related reading

- Shadow AI agents: how to find the agents nobody registered
- What is AI agent identity?
- What is non-human identity (NHI)?
- Compliance mapping: NIST, CISA, and the agentic frameworks

References (public)

- Okta, "Okta for AI Agents" product page (fetched June 2026): https://www.okta.com/products/govern-ai-agent-identity/
- Microsoft Learn, "What is Microsoft Entra Agent ID?" (fetched June 2026): https://learn.microsoft.com/en-us/entra/agent-id/what-is-microsoft-entra-agent-id
- Astrix Security, platform product page (fetched June 2026): https://astrix.security/product/
- Oasis Security, "Agentic Access Management" product page (fetched June 2026): https://www.oasis.security/agentic-access-management
- Gartner press release, "Gartner Predicts 40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026" (August 26, 2025): https://www.gartner.com/en/newsroom/press-releases/2025-08-26-gartner-predicts-40-percent-of-enterprise-apps-will-feature-task-specific-ai-agents-by-2026-up-from-less-than-5-percent-in-2025
- OWASP GenAI Security Project, "OWASP Top 10 for Agentic Applications" (December 2025): https://genai.owasp.org/
- NIST NCCoE, "Accelerating the Adoption of Software and Artificial Intelligence Agent Identity and Authorization" concept paper (February 2026): https://csrc.nist.gov/pubs/other/2026/02/05/accelerating-the-adoption-of-software-and-ai-agent/ipd