SPECTEROPS

AI Security starts with eliminating attack paths

As you're reading this, AI-enabled attackers have millions, if not billions, of paths to your critical assets. They can move from foothold to impact in minutes, exploiting the access that your environment already grants. Agent governance is important, but the urgent priority should be eliminating attack paths.

SPECTEROPS.IO

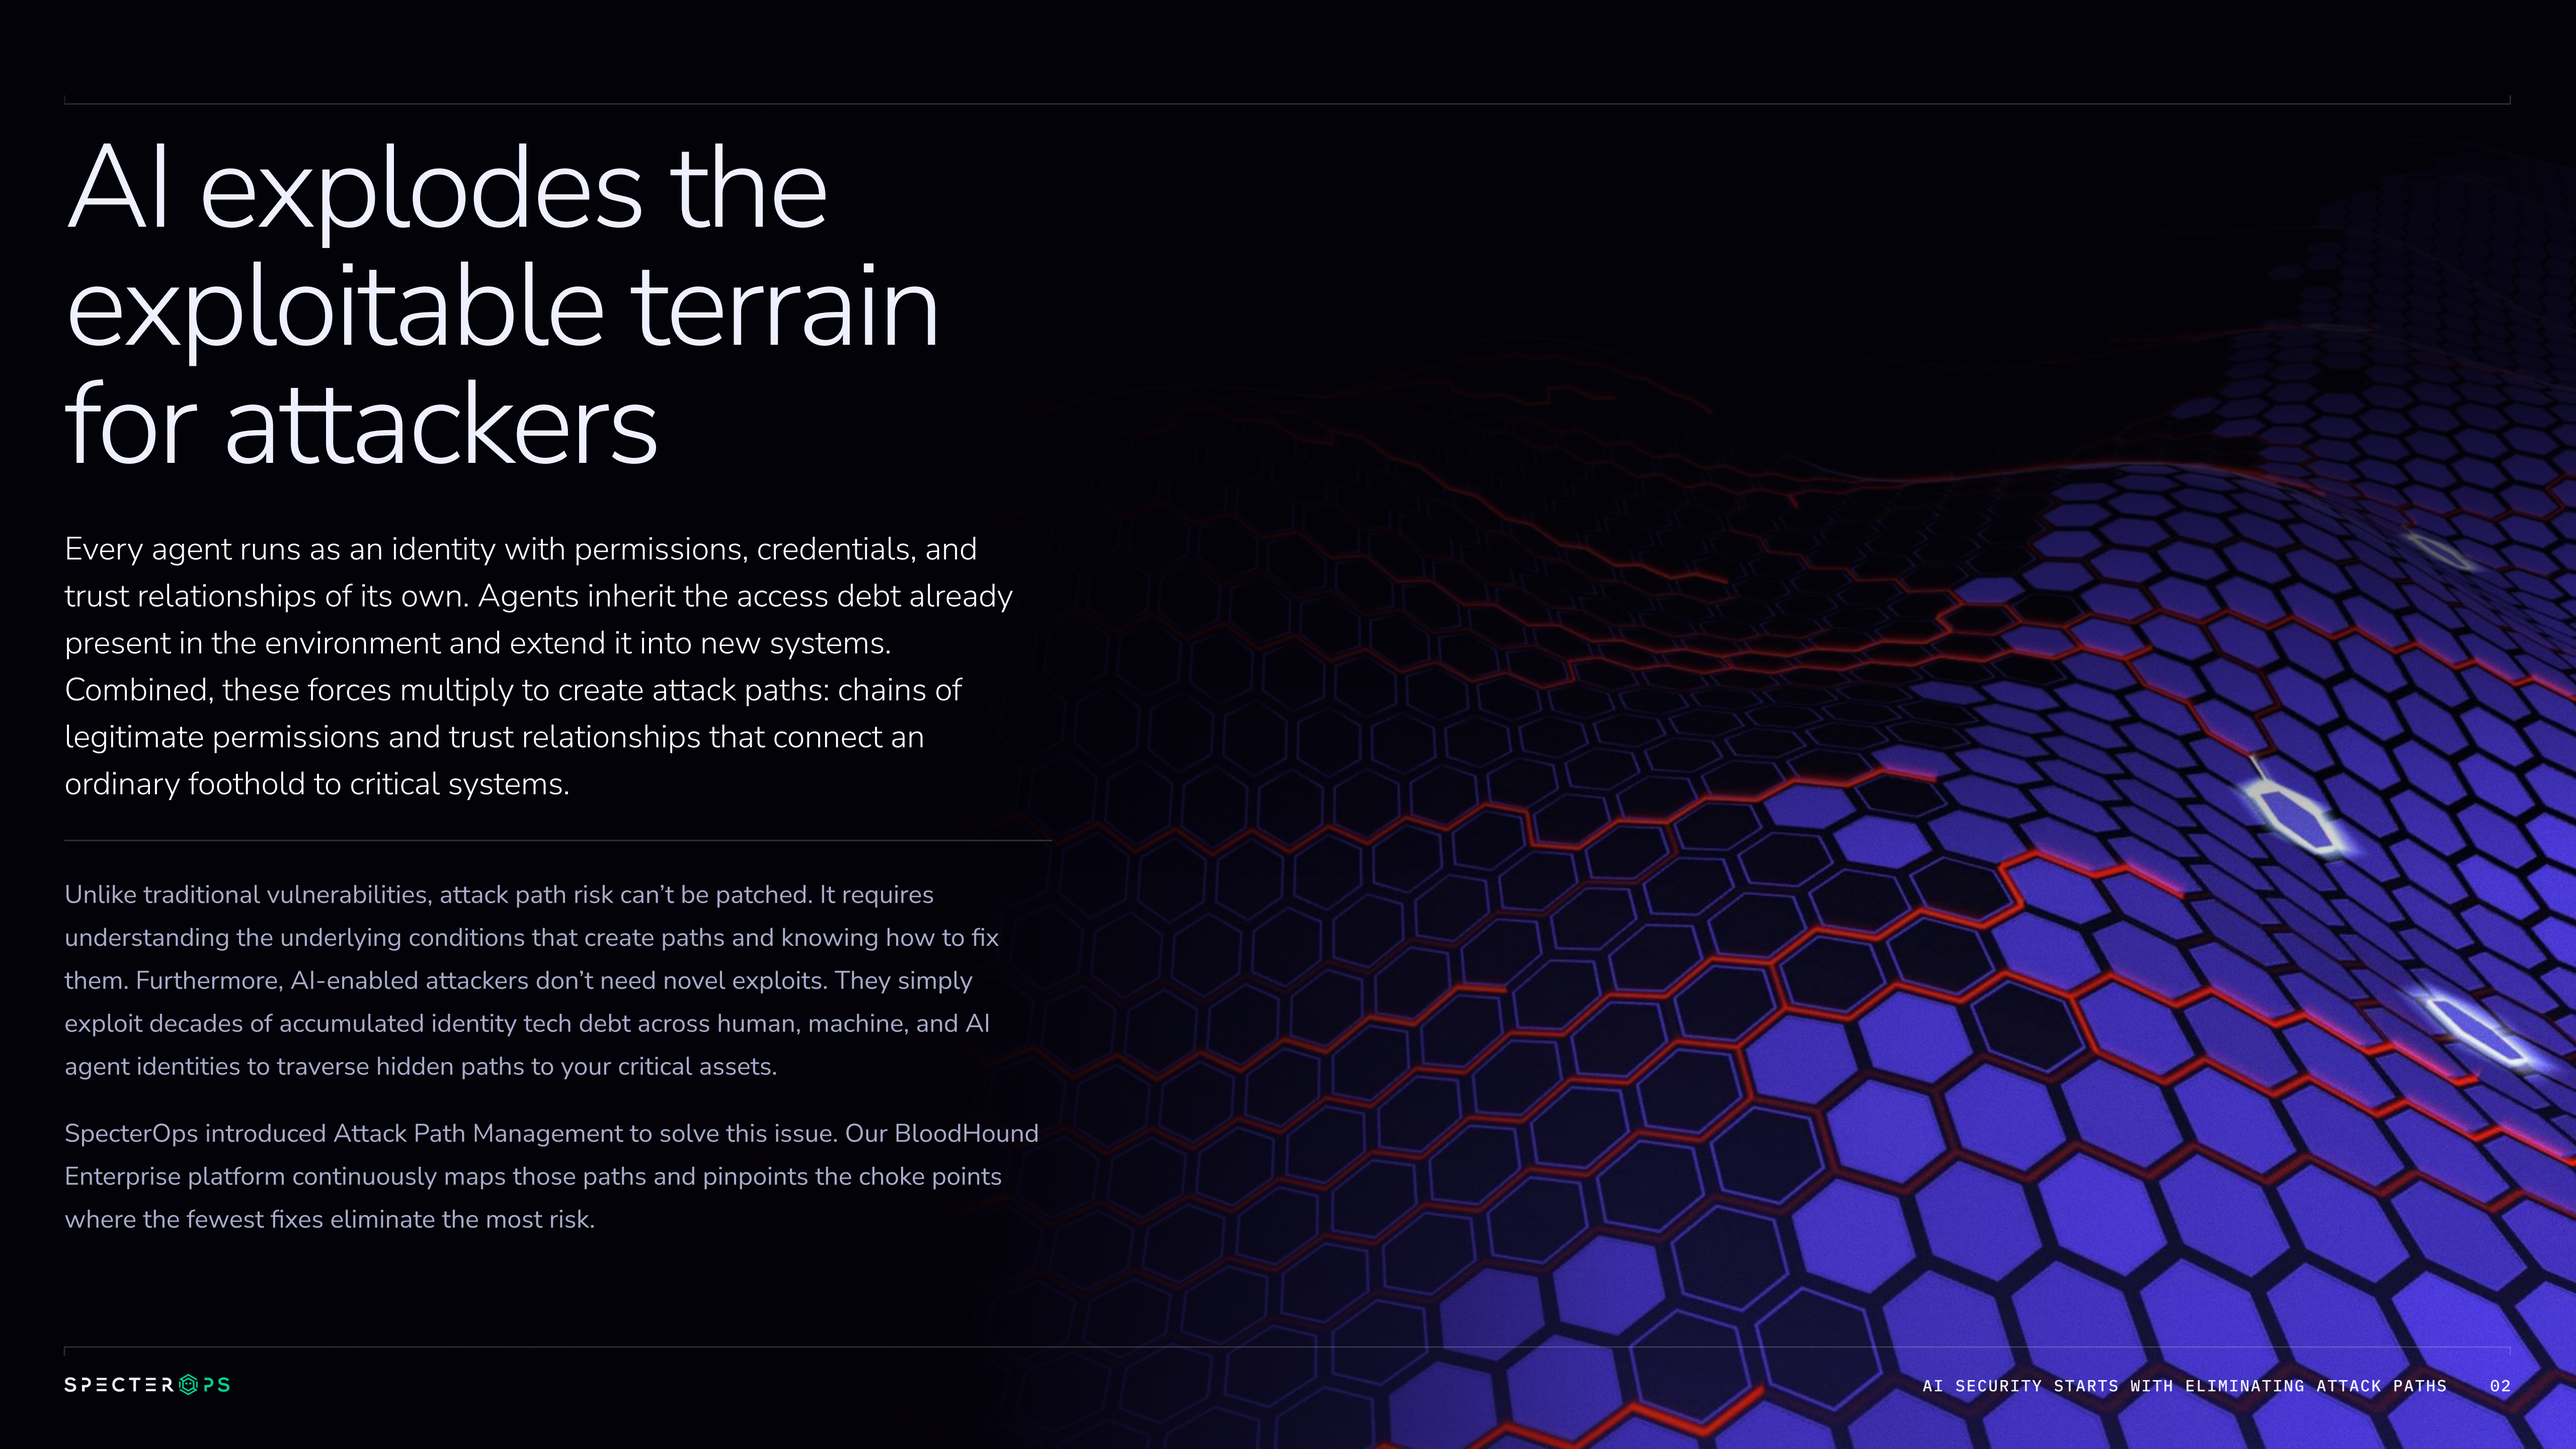


AI explodes the exploitable terrain for attackers

Every agent runs as an identity with permissions, credentials, and trust relationships of its own. Agents inherit the access debt already present in the environment and extend it into new systems. Combined, these forces multiply to create attack paths: chains of legitimate permissions and trust relationships that connect an ordinary foothold to critical systems.

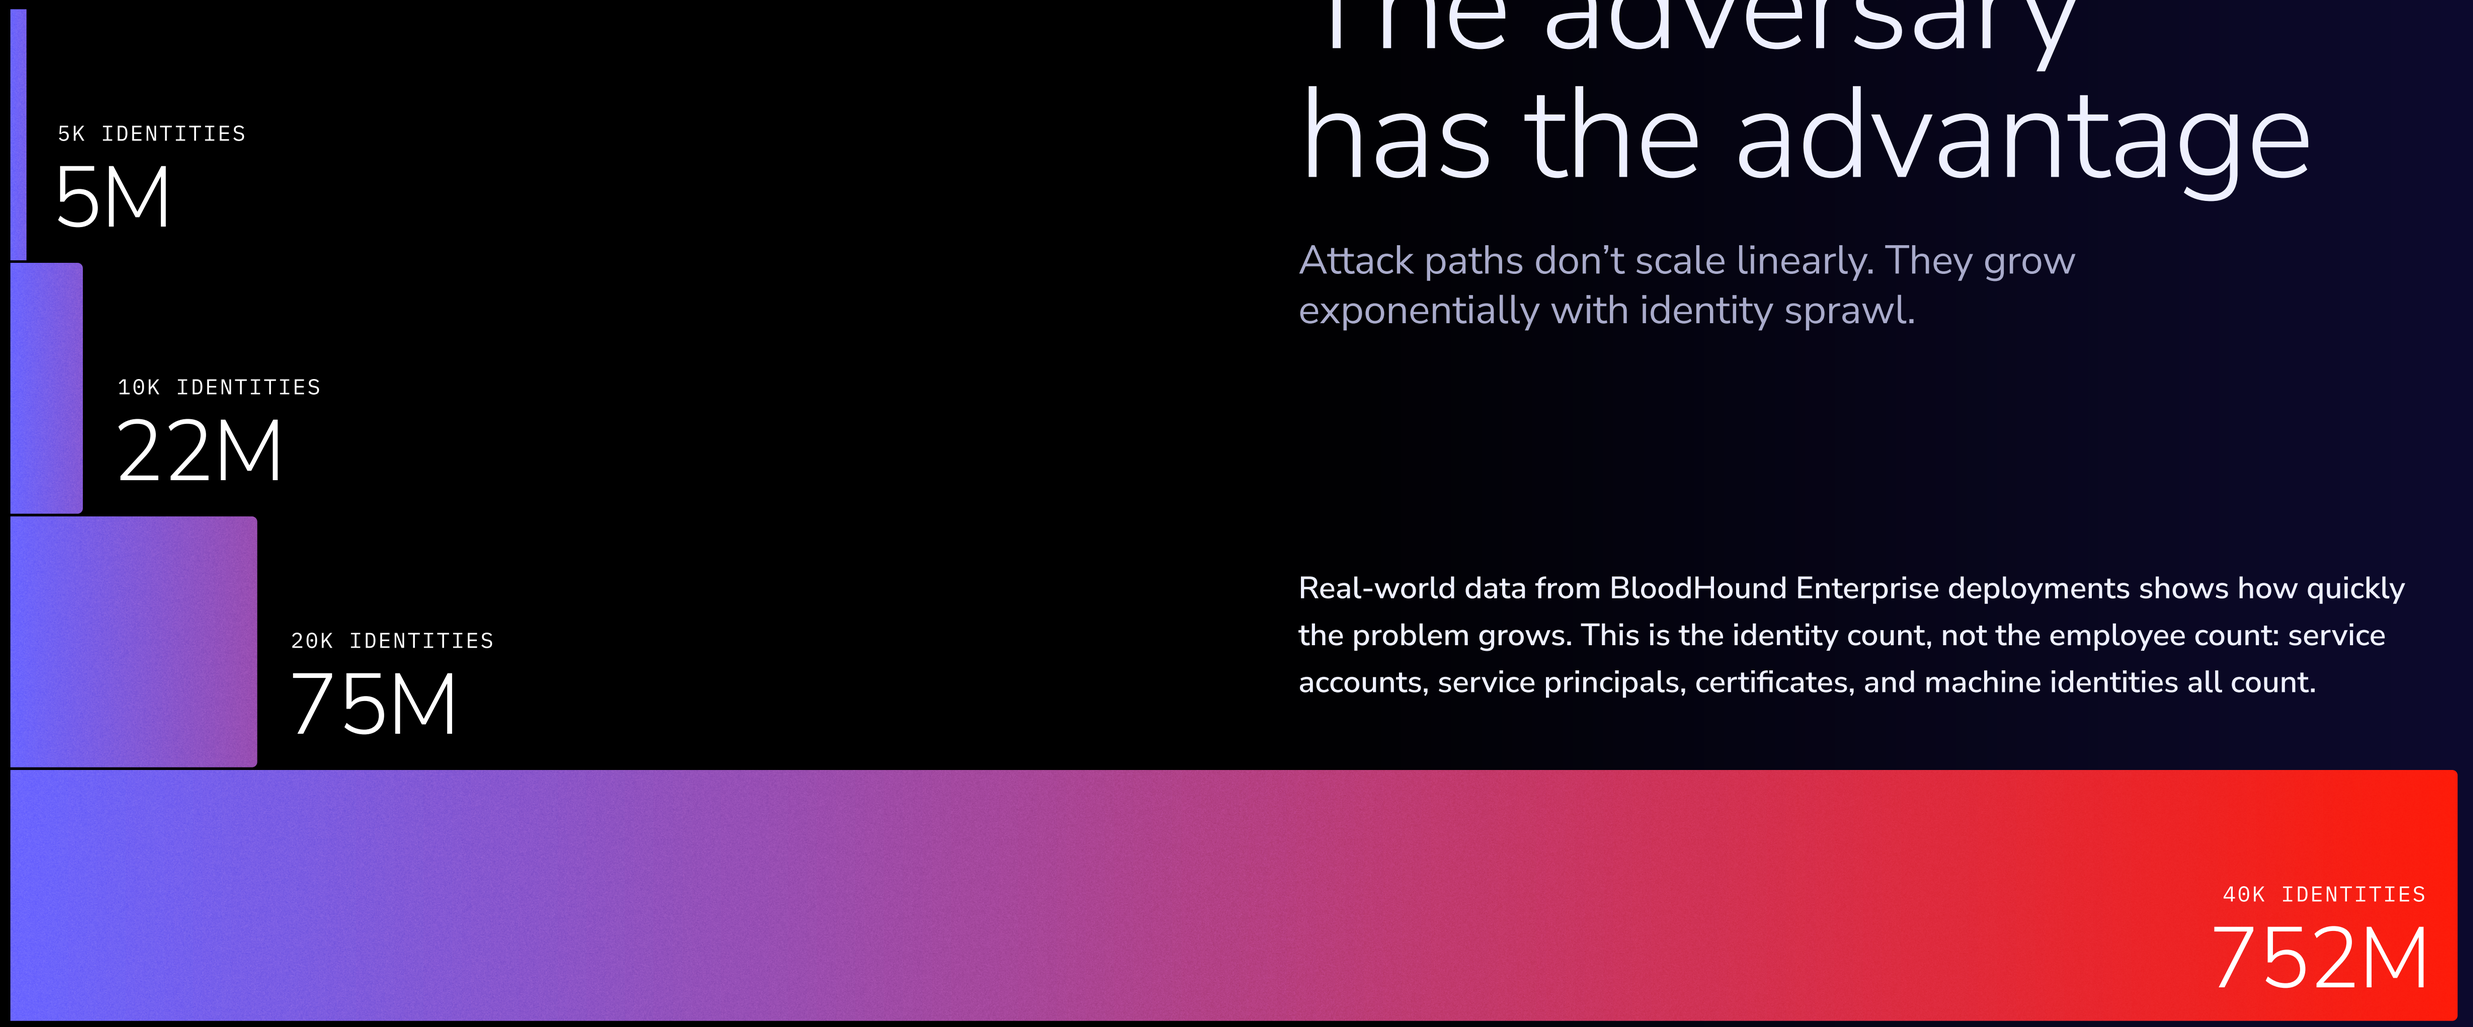
Unlike traditional vulnerabilities, attack path risk can't be patched. It requires understanding the underlying conditions that create paths and knowing how to fix them. Furthermore, AI-enabled attackers don't need novel exploits. They simply exploit decades of accumulated identity tech debt across human, machine, and AI agent identities to traverse hidden paths to your critical assets.

SpecterOps introduced Attack Path Management to solve this issue. Our BloodHound Enterprise platform continuously maps those paths and pinpoints the choke points where the fewest fixes eliminate the most risk.



The adversary has the advantage

Attack paths don’t scale linearly. They grow exponentially with identity sprawl.



THE SCALE OF ATTACK PATHS GROWS EXPONENTIALLY AS AN ORGANIZATION’S IDENTITIES INCREASE.

Source: BloodHound Enterprise deployment data (State of Attack Path Management)

Real-world data from BloodHound Enterprise deployments shows how quickly the problem grows. This is the identity count, not the employee count: service accounts, service principals, certificates, and machine identities all count.

AI turns tech debt into escalation paths on demand

Mapping the obscure corners of an environment used to be slow, manual work. Attackers had to decide where to spend limited operator time.

01 Unfamiliar systems needed investigation

02 Alternate routes took time to explore

03 Custom tooling needed development and testing.

Those constraints forced attackers to abandon some opportunities. Defenders benefited from the delay.

AI removes much of that friction for attackers. It analyzes configurations, builds tooling for what it finds, and runs investigations in parallel. Legacy identity decisions and misconfigurations that once sat too deep in the environment to justify the effort now take minutes to find and use.

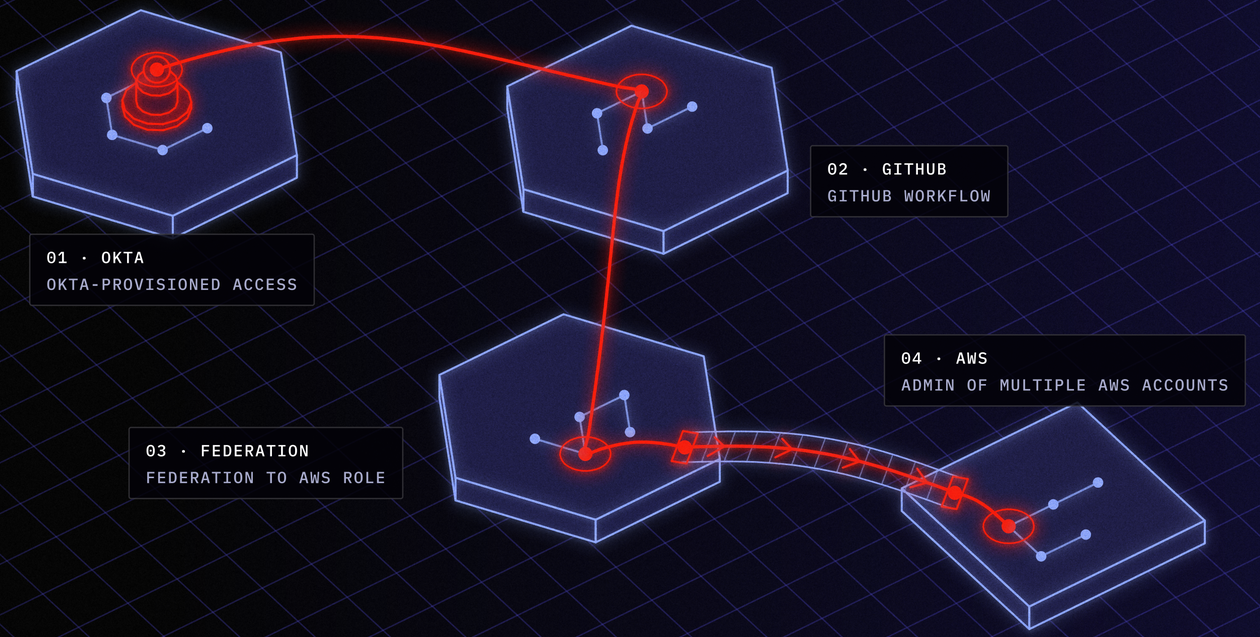
AI Attackers move freely between these identities and systems

BLOODHOUND ENTERPRISE

During a recent SpecterOps engagement, BloodHound Enterprise identified a path to administrative control of multiple AWS accounts. The path began with Okta-provisioned access to GitHub. From GitHub, workflow and federation configurations let a deployment task assume overprivileged AWS roles. The task existed to publish an application to an S3 bucket. In practice, it handed over administrative control of multiple AWS accounts.

Each configuration looked reasonable in isolation. Only cross-platform analysis showed what they enabled together.

Any AI agent able to influence that workflow would inherit the same control. Narrowing the trusted workflow contexts and role permissions breaks the chain.



Your AI agents inherit the access you already granted

> DATA SPOTLIGHT

88%

of large-enterprise security leaders already use AI agents or plan to within the next 12 to 18 months.

Source: SpecterOps / Omdia [2]

Enterprises are putting AI agents to work quickly. Eighty-eight percent of large-enterprise security leaders already use AI agents or plan to within the next 12 to 18 months [2]. More than half of CISOs (55%) already run MCP servers or other agent-to-application integrations in production [3]. These deployments inherit years of identity tech debt that survived security assessments, governance reviews, and modernization efforts.

Each agent runs as an identity. It holds credentials, carries delegated permissions, and connects to the systems it was built to work with. An agent's approved task may be limited, but its underlying permissions let a compromised agent or someone controlling it reach far beyond that purpose.

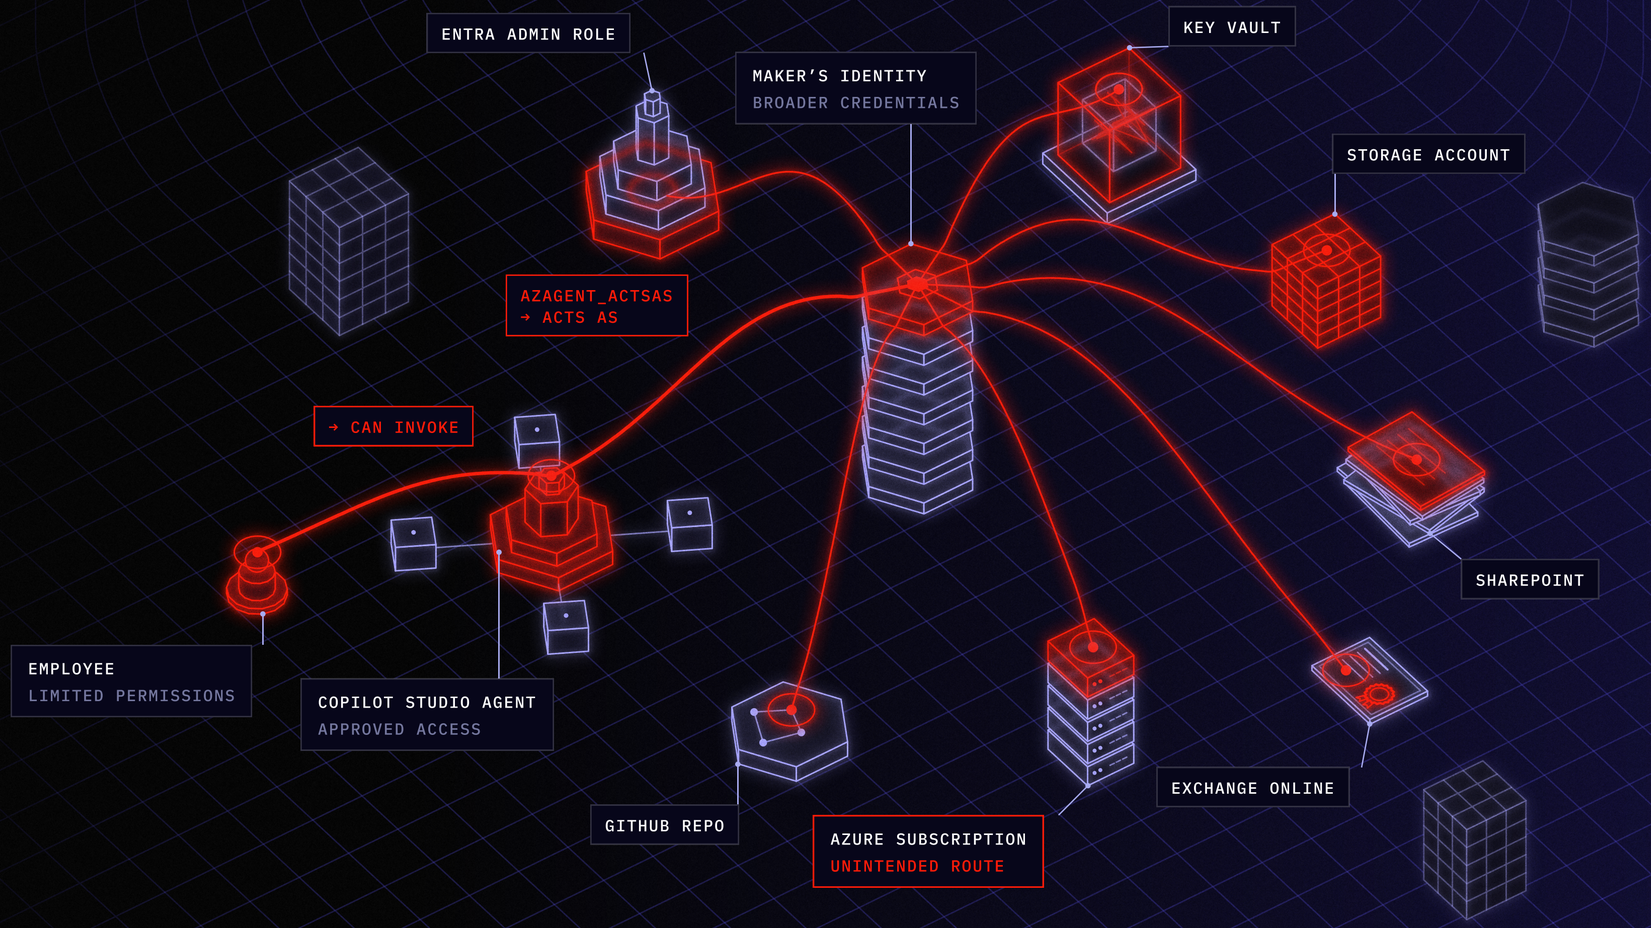
[2] SpecterOps and Omdia, 2026 Trends in Identity Attack Path Management. Survey of 541 cybersecurity decision-makers at organizations with 10,000+ employees, January 2026.

[3] Steph Zhang and Owen Colby, "The New Security Agenda," a16z, October 28, 2026. Survey of 100 CISOs at companies with 1,000+ employees.

Human and agentic identities both pass when reviewed separately

Consider an employee with limited permissions and approved access to a Copilot Studio agent. The agent uses its maker's broader credentials to reach SharePoint and Azure resources the employee can't access directly. Review the employee's entitlements and the agent's separately, and both pass. Map the connections and the unintended route becomes visible.

Attack Path Management shows which inherited permissions let an agent reach critical assets, and where one change removes the most risk. Fix those choke points, and every agent deployed afterward inherits a safer environment.



“Conventional wisdom for managing AI agent identity is to provide Joe’s agent with Joe’s access and permissions. Unfortunately, that approach doesn’t consider the existing attack paths in the environment, it just amplifies the problem. It means now wherever Joe’s access is vulnerable, so is the agent’s.”

JARED ATKINSON, CTO, SPECTEROPS

The number one priority for CISOs should be cutting attack paths

> DATA SPOTLIGHT

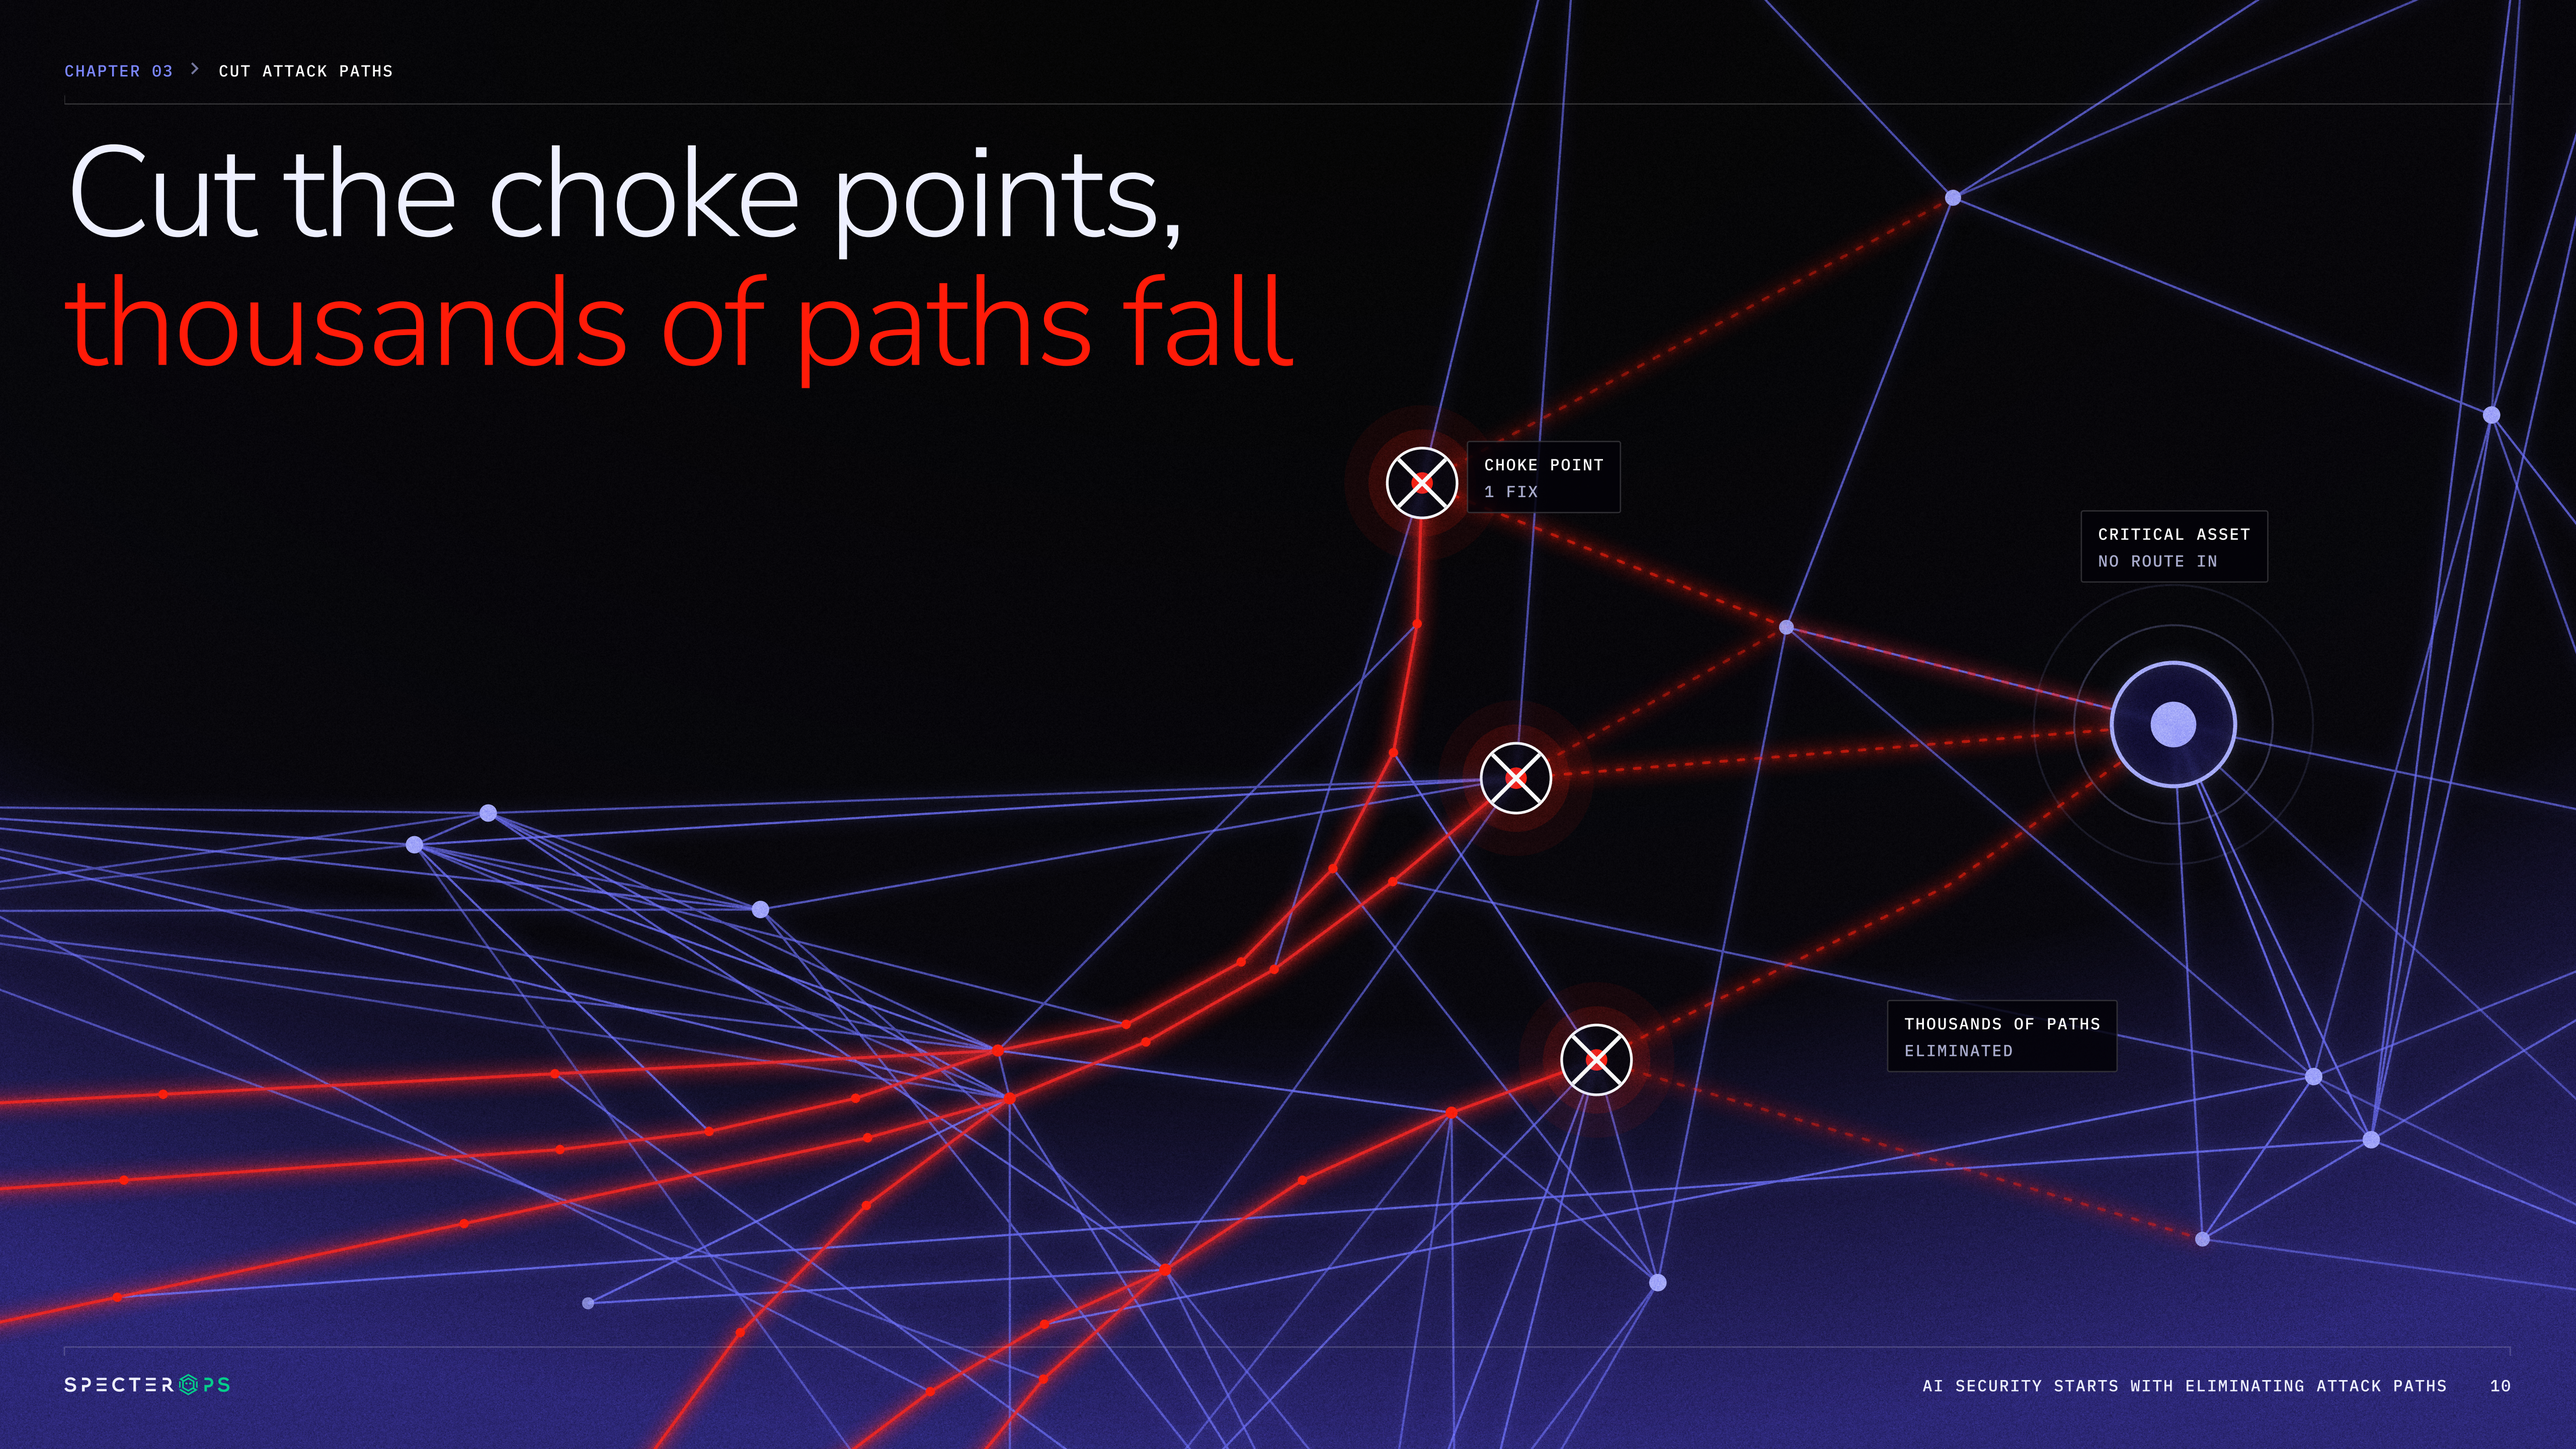
1B

attack paths on average across SpecterOps customer environments.

Source: SpecterOps platform data

Vulnerability management and agent governance are foundational to AI security, but the urgent challenge is identifying and severing attack paths. An organization can patch every known vulnerability and pass every access review while an attacker holds a working route to its critical systems. Those routes run through legitimate permissions and trust relationships. Patching never touches them, and no single entitlement review reveals them.

Cut the choke points, thousands of paths fall



Where traditional identity security falls short

Identity governance reviews entitlements one at a time

Governance tools ask whether each identity should hold each entitlement. Attack paths form where entitlements combine across agents, workflows, and connected systems. An agent that runs on its maker's credentials, a GitHub workflow that an AWS role trusts, a help desk group that can reset an administrator's password: each can pass review on its own. Identity governance manages who should have access. It has no view of what those permissions enable together.

Finding attack paths leaves the hardest work undone

Enterprise environments can contain millions of attack paths. Discovery at scale proves an environment is exposed, but it offers little guidance on which change will reduce the most risk. Closing the path shown in a demo or an offensive security test report can leave equivalent access through thousands of others, and 36% of security decision-makers say they lack the trained staff to validate and fix identity-based findings [2].

A persistent human attacker might take days or weeks to find those alternate routes. A swarm of AI agents working in parallel can find them instantly.

[2] SpecterOps and Omdia, 2026 Trends in Identity Attack Path Management. Survey of 541 cybersecurity decision-makers at organizations with 10,000+ employees, January 2026.

Every control sees part of the path, none sees the whole route

Enterprise security controls address important parts of identity risk, but attack paths connect permissions and trust relationships across their boundaries.

ITDR · IDENTITY THREAT DETECTION AND RESPONSE

An attacker using a compromised account's authorized permissions can perform actions that appear acceptable to Identity Threat Detection and Response (ITDR).

PAM · PRIVILEGED ACCESS MANAGEMENT

Stolen credentials or session tokens can enable access outside workflows governed by Privileged Access Management (PAM).

IGA · IDENTITY GOVERNANCE

Identity governance can verify who is entitled to access a critical system while overlooking who could take over an entitled account. Each control provides value, yet the combined route from an ordinary foothold to privileged access can remain unseen.

CTEM · CONTINUOUS THREAT EXPOSURE MANAGEMENT

Continuous Threat Exposure Management (CTEM) programs can identify and prioritize individual exposures, but they don't provide a connected view of how identity relationships turn those exposures into routes to critical assets. Enterprise AI adoption extends those relationships further. Agents receive credentials and delegated permissions to perform work, and their reach can extend beyond their assigned task through the accounts, systems, and trust relationships they can control.

Attack Path Management addresses the gaps across these tools, providing a continuous view and helping teams remove the underlying access debt, strengthen the boundaries their existing controls support, and reduce the routes available to an attacker.

AI-centric security just adds another silo

An army of AI agents can accelerate offensive testing, but the number of attacks executed or paths discovered says little about how much exposure remains. Years of accumulated permissions and trust relationships create millions of potential routes to critical assets, with many sharing the same overprivileged accounts and delegated access. Fixing demonstrated routes can leave alternatives intact when the underlying privileges remain. Enterprise AI adoption compounds that debt as agents inherit existing access and introduce new privileges, giving attackers more opportunities to move through legitimate permissions.

01 **Eliminate paths at scale:** APM identifies shared chokepoints where a single change can remove thousands of attack paths.

02 **Prioritize the greatest impact:** Focuses security and identity teams on the changes that remove the most exposure to critical assets.

03 **Address paths beyond those demonstrated:** Maps the underlying identity relationships that make alternative routes possible.

04 **Sustained improvement:** Continuously tracks exposure as permissions, trust relationships, and AI identities change.

SPECTEROPS

Leave attackers nowhere to go

After a decade of simulating advanced adversaries, SpecterOps knows an attacker's worst outcome is to gain access to an environment and find they have nowhere to go. When AI can search every route, defenders need to remove the routes that matter first.

Attack Path Management works from the other end. It starts with critical assets and traces the direct and indirect access that makes them reachable across every connected system. BloodHound Enterprise runs that analysis continuously across on-premises, cloud, and hybrid environments, including the identities and permissions AI agents introduce. It finds the choke points where one change removes thousands of paths and guides the fix. Because the analysis starts from what an attacker could reach, it holds up even when the entry point is unknown.

01 See what adversaries see. Map identity relationships across Entra ID, Okta, AWS, GitHub, Active Directory, and hybrid environments in a single attack graph.

02 Follow every agent. Trace who can invoke or modify an AI agent, whose identity it runs as, and what that identity can reach across Entra Agent ID, Copilot Studio, and Power Automate.

03 Fix the choke point. Prioritize the changes that remove the most attack paths to critical assets, with guidance on how to make them.

04 Catch new exposure as it appears. Re-analyze continuously as new identities, integrations, and agents come online.

OpenAI ANTHROPIC AISI AI SECURITY INSTITUTE

BLOODHOUND ENTERPRISE CUSTOMER RESULTS

2 *trillion*

attack paths closed by BloodHound Enterprise customers

SPECTEROPS PLATFORM DATA

54%

reduction in attack paths within six months

SPECTEROPS PLATFORM DATA

90%

reduction in attack paths within 12 months

SPECTEROPS PLATFORM DATA



Key takeaways

For security leaders

01 Your environment has more than 1 billion attack paths to your most critical assets. These paths cross platform boundaries, chaining privileges from one system into the next, such as Okta to GitHub to AWS.

02 Prioritize eliminating the paths with direct access to your critical systems and reduce any exposure AI may introduce. By addressing strategic chokepoints, you spend scarce time and effort where it matters most.

03 Measure progress by the number of attack paths removed from your environment, giving your board a clear picture of risk reduction.

For practitioners

01 An AI agent is just another identity, but it amplifies the Attack Path Management problem. Know who can invoke it, whose credentials it uses, and what those credentials reach.

02 Prioritize choke points over individual findings. The right fix can remove thousands of paths.

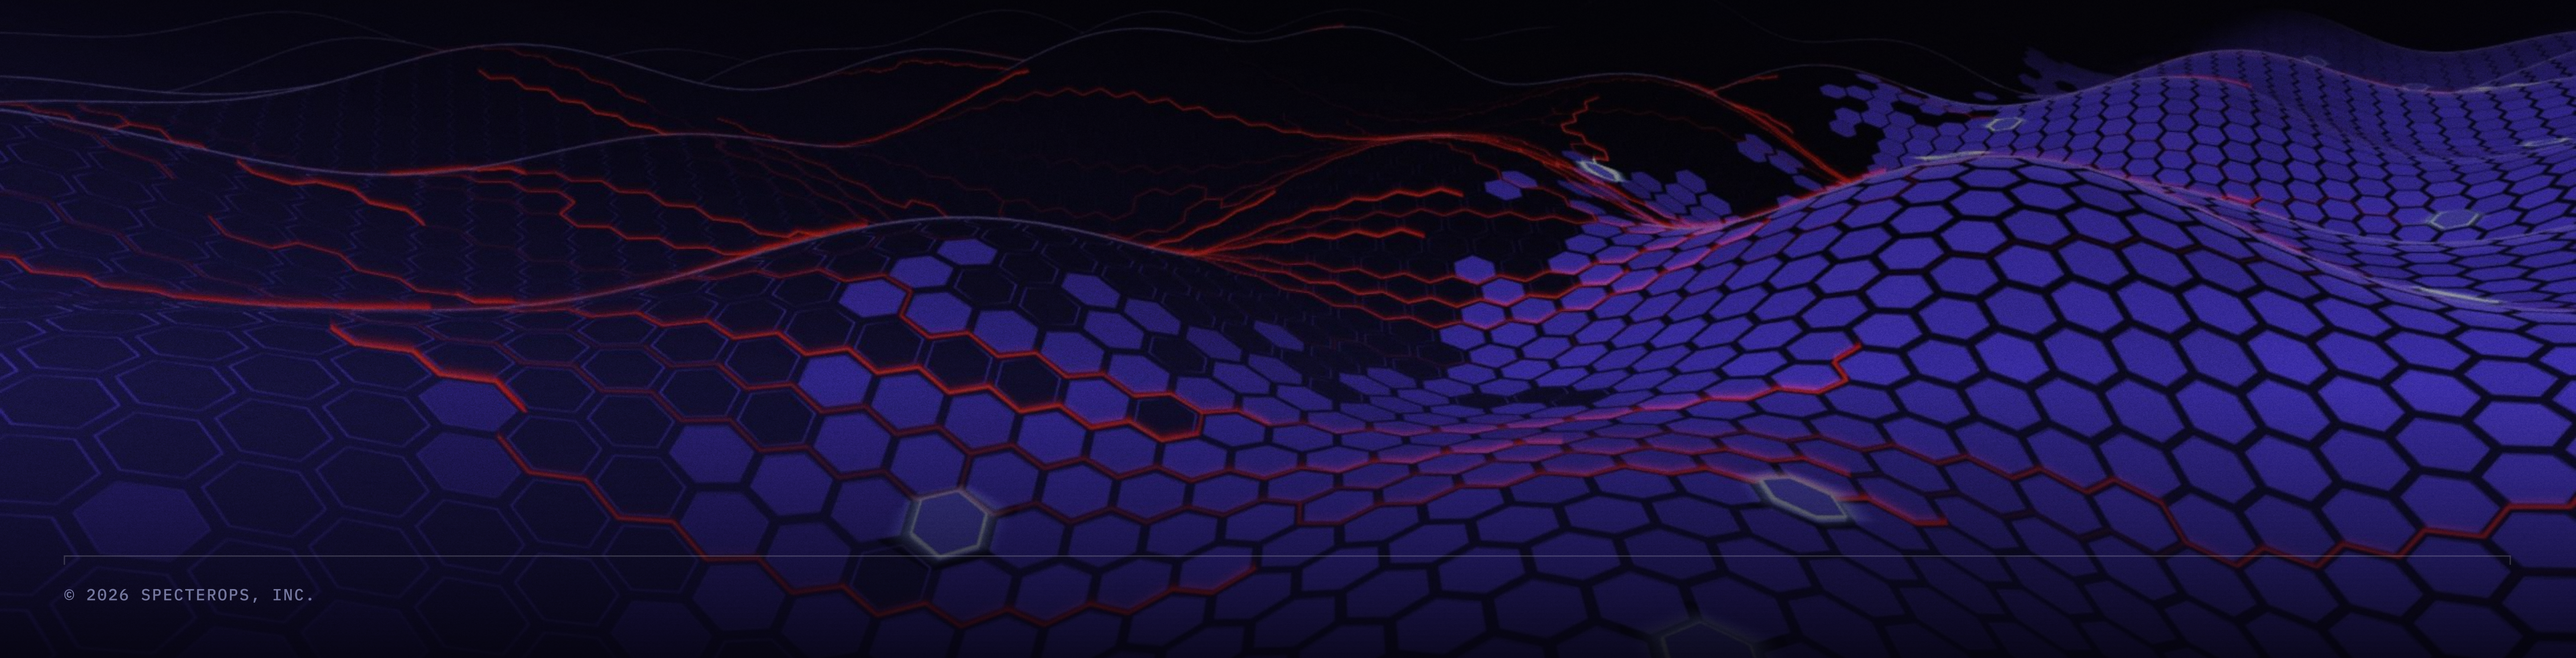
03 Don't limit visibility to just one identity provider. Ensure coverage of all your interconnected platforms, such as Okta, GitHub, and AWS. These identities are often interconnected, meaning risks can occur on either side and impact your entire estate.

SPECTEROPS

Know the AI adversary

Build the AI-enabled defense your environment needs to leave attackers nowhere to go. Contact us at specterops.io/contact.

Learn more at specterops.io ›



© 2026 SPECTEROPS, INC.