



DATASHEET

POST-AUTHENTICATION IDENTITY GOVERNANCE

THE ATTACK ALREADY GOT PAST YOUR MFA.

87%

of 2024 breaches involved identity exploitation after valid authentication. Session hijacking, token replay, and lateral movement through trusted relationships are post-authentication problems. MFA was never designed to stop them. – Unit 42 Incident Response Report, 2025

Non-human identities and AI agents are the fastest-growing ungoverned attack surface in the enterprise. They hold privileged access, operate silently, and carry no lifecycle controls — making them the preferred path for lateral movement after a valid login.

Agent-Based

AI Agents & NHI · MCP Server Governance

Zero Standing Privileges



WHITESWAN
IDENTITY SECURITY

The Problem

LATERAL MOVEMENT HAPPENS AFTER THE LOGIN

Attackers don't need to break down the door. They walk through it with a valid session, then chain identity exploits across your environment – endpoint to domain controller to cloud to AI agents. Every step exploits an identity your current stack doesn't govern.

HOW THE ATTACK CHAIN UNFOLDS - AND WHERE WHITESWAN INTERCEPTS EACH STEP

ENTRY POINT

MFA Bypassed via Session Hijacking or Token Replay

A valid authentication completes. The attacker steals the session token – not the password. They're inside with a legitimate identity. No alert fires. MFA has done its job. The attack has not yet started in the way your tools are designed to detect it.

LATERAL MOVE 1

Endpoint Privilege Escalation

Local admin accounts and service accounts on the compromised endpoint provide the first escalation path. These identities were never onboarded into your PAM vault – they carry standing privileges your tools have no visibility into.

 **WS Identity Vault — discovers and governs every local and service account**

LATERAL MOVE 2

Active Directory Exploitation — DCSync, Pass-the-Hash, Golden Ticket

Using extracted credentials, the attacker reaches a domain controller and begins replicating password hashes or forging Kerberos tickets. From here, every system in the domain is accessible. This move is fast and largely invisible to tools that don't sit on the DC itself.

 **WS Identity Vault AD ITDR — edge processing on every DC detects this in real time**

LATERAL MOVE 3

Cloud IAM & Pipeline Credential Abuse

With domain access, CI/CD pipeline tokens, AWS IAM roles, and Kubernetes service accounts become accessible. Most carry standing admin privileges with no expiry. The attacker moves from on-prem to cloud with credentials that have been sitting dormant – completely valid, never monitored.

 **WS Cloud Mesh — JIT access eliminates standing cloud credentials across every provider**

LATERAL MOVE 4

AI Agent & MCP Server Access — the Surface No Tool Governs Whole

AI agents already have standing admin keys to production databases, CRM systems, and code repositories. MCP servers bridge AI models directly into core infrastructure. No PAM tool governs them. No EDR tracks them. They are the fastest-growing ungoverned surface in the enterprise – governed in fragments, by tools that don't see the rest of the chain.

 **WS Agentic Gateway — full lifecycle governance for every AI agent and MCP server**

WHY THE CURRENT STACK FAILS

Fragmented by design

Cloud JIT tool. Standalone NHI product. ITDR solution. SIEM. Endpoint privilege manager. Legacy PAM vault. Each governs one piece – none govern the connections between them. That's exactly where attackers move.

No single tool governs AI/NHI across the chain

Traditional PAM and IAM were built for human users. API keys, pipeline tokens, AI agent credentials, and MCP server access are invisible to every tool in the current stack. They carry admin-level access with no lifecycle controls.

The newest identity tools are built cloud-first – which leaves the legacy on-prem and Active Directory estate, where the oldest standing privileges live, out of scope. Whiteswan governs that estate, the cloud, and the agentic surface in the same engine.

Platform Capabilities

Non-Human Identities & AI Governance

Whiteswan governs every non-human identity across all three attack surfaces – on-prem service accounts, cloud workload identities, and AI agents – through one unified policy engine. Most tools govern one surface and call it coverage. But lateral movement doesn’t respect surface boundaries – it crosses them. A chain governed in three places, by three tools, with three blind spots between them, is not governed. Whiteswan closes it end to end, from one engine.

Identity Type	Security Challenge	Risk	Whiteswan Solution
Service & Privileged Accounts	High system access, rarely rotated, not in PAM vault	Critical	Legacy & On-Prem PAM · AD ITDR
API Tokens & Pipeline Credentials	Widely distributed, hardcoded, no expiry	Critical	Cloud Entitlements · JIT Access
Cloud Workloads & K8s Accounts	Created faster than governance can track	High	Multi-Cloud Policy Engine · Auto-Revocation
AI Agents (Claude, GPT, Custom)	Admin-level standing keys to production systems	Critical	WS Agentic Gateway · JIT Agent Policies
MCP Servers (AI Control Plane)	Bridges AI models into core infrastructure	Critical	MCP Server Lifecycle Governance
RPA Bots & Automation Tokens	Automate access to critical systems, no controls	High	NHI Management · Credential Rotation

WS Agentic Gateway — AI Agents & Non-Human Identities

The Governance Platform Built Across Your Whole Identity Estate.

The enterprise-grade security control plane for AI agents and MCP server ecosystems — govern, inspect, and protect every tool interaction at scale.

AI Agent Governance

JIT policies for every AI agent accessing CRM, code repositories, and production databases. No standing admin keys for bots – ever. Every agent call is inspected, authorized, and governed before execution.

MCP Server Lifecycle

An end-to-end MCP governance lifecycle – the AI control plane bridging AI models into production – Discover, Govern, JIT, Monitor, Revoke – enforced at the protocol chokepoint.

Capability Registry & Tool Control

Controls which tools AI agents can invoke. Blocks untrusted tool execution. Only validated, approved capabilities are exposed to agents – eliminating supply-chain attacks via malicious MCP skills.

Policy & Authorisation Engine


RBAC, argument validation, and scoped access enforced at the gateway layer. Every tool call passes through policy enforcement before it reaches any enterprise system. No exceptions.

Data Flow Guardrails

Detects and blocks unsafe data exfiltration patterns. Prevents cross-tool data leaks where attackers chain multiple tool calls to extract sensitive information across AI agent interactions.

NHI Management — API Keys & Tokens

API keys, pipeline credentials, and automation tokens governed end-to-end. Continuous discovery flags orphaned and over-privileged machine identities. No forgotten secrets, no standing access.



Platform Capabilities

Conventional Tools Vs Whiteswan

Traditional PAM and IAM platforms were built before AI agents existed. They govern human sessions – not machine identities, not AI agents, not MCP servers. Every vendor in the current landscape covers a slice of the problem. Whiteswan closes the chain between the slices – from legacy on-prem service accounts through multi-cloud workload identities to AI agent ecosystems – in a single unified policy engine.

Capability	Whiteswan	Traditional IAM/PAM	CSP Native Tools
NHI Auto-Discovery & Classification	✓	✗	Partial
AI Agent Governance (JIT per call)	✓	✗	✗
MCP Server Lifecycle Governance	✓	✗	✗
In-Flight Tool Call Interception	✓	✗	✗
Lateral Movement Detection (AD ITDR)	✓	Partial	✗
Multi-Cloud NHI Governance	✓	✗	Partial
Legacy OS Coverage (Win 2003–2022, RHEL 6–9)	✓	Partial	✗
Operates Without a Sensor Dependency	✓	Partial	✗
Deployment Time	30 Days	Weeks – Months	Varies

IMPLEMENTATION & DEPLOYMENT



30-Day Pilot

JIT access and monitoring live in 30 days. No infrastructure overhaul required. Production-ready governance from day one without displacing existing tooling.



Agent-Based Primary

Lightweight agents on Windows, Linux, & macOS. Optional agentless gateways for constrained environments. Additive to every existing security investment – it works alongside your stack, no sensor required.



Hybrid Deployment

On-prem, cloud, and AI agent governance from one unified policy engine. Full support for AWS · Azure · GCP across all workload identity types.



Native Integrations

CrowdStrike · Splunk · XDR platforms. Whiteswan is additive to every deal – enhancing, not replacing, your existing security stack.



WHITESWAN

IDENTITY SECURITY

Non-human identities and AI agents were never onboarded, never governed, and never revoked. Whiteswan closes that gap across on-prem, cloud, and agentic surfaces – governed from a single policy engine, additive to the stack you already run, without displacing what you have.

**POINT TOOLS GOVERN A SURFACE.
WHITESWAN GOVERNS THE CHAIN.**

API keys, pipeline credentials, AI agents, and MCP servers – are ungoverned in every enterprise. Whiteswan closes that surface across on-prem, cloud, and AI environments from a single policy engine.

- **NHI Risk Assessment:** Discover every ungoverned NHI and AI agent – including the ones your existing tools cannot see.
- **30-Day Pilot:** JIT access, continuous monitoring, and automatic revocation live in 30 days, without infrastructure changes.
- **ROI Analysis:** Calculate the risk reduction value of governing every non-human identity in your environment.



vmamidi@whiteswansecurity.com



www.whiteswansecurity.com