DAYBLINK CONSULTING

Securing Non-Human Identities in the AI Era

How DayBlink Consulting modernized a Non-Human Identity program at a Fortune 500 bank to meet the demands of agentic AI

Introduction

A large bank was concerned that its identity program was not keeping pace as it accelerated adoption of AI agents, automation pipelines, and local LLMs. Its population of non-human identities (NHIs) such as service accounts, workload identities and agents was growing faster than the traditional IAM models built to govern them. Every new agent, MCP server, and CI/CD workflow introduced credentials that must be inventoried, scoped, rotated, and eventually retired.

For large financial institutions, the stakes are amplified. Regulators expect demonstrable control over privileged access, yet many of the fastest growing identity types now originate at the edges of the enterprise: in developer tooling, agentic workflows, and locally run automation, often outside centralized IAM visibility.

To get ahead of this shift, the DayBlink Consulting team evaluated the NHI management program against the emerging frontier AI threat landscape, and delivered a defensible roadmap, grounded in evidence, for the next two years of NHI investment.

Problem

Our client's strong IAM foundations were built before the rise of agentic AI, and leadership sought the better understand where and if this new identity risk required additional investment

The bank's IAM organization had established many identity best practices that it was extending to NHI, including: inventory linkage, lifecycle controls, and Zero Trust principles, but the landscape was changing rapidly. AI agents, automation, and local execution models were materially increasing the number, autonomy, and privilege of non-human identities across the enterprise, frequently outside centralized IAM visibility and control.

The greatest exposure sat at the edges. Local MCP servers, developer tooling, CI/CD automation, and agentic workflows introduced identities that could bypass traditional governance, rely on locally stored credentials, and operate beyond existing detection and monitoring models. Meanwhile, governance and control patterns for agentic AI remained immature across the industry, providing no established playbook.

Leadership faced a prioritization challenge. Without a structured, empirical view of current NHI maturity, it was difficult to determine which gaps were acceptable, which demanded investment in process, tooling, or capacity, and how to sequence that investment credibly for executives, auditors, and regulators.

Solution

DayBlink Consulting built an empirically grounded NHI inventory, evaluated identity maturity, updated the team's operating model, and then translated it into a prioritized two year roadmap

DayBlink Consulting partnered with the client to evaluate its NHI program through a structured analysis across three complementary methods, designed to ensure defensibility, traceability, and alignment to the scope owned by IAM.

First, we conducted discovery sessions with IAM leaders and technical experts across the cyber organization (spanning identity governance, secrets and credential management, PAM, cloud IAM and workload identity, and AI agent architecture and strategy) to validate operating model realities, tooling coverage, lifecycle friction points, and ownership boundaries.

Second, we performed a maturity assessment across five NHI management domains owned by IAM: Governance & Ownership, Inventory & Classification, Authentication & Credential Controls, Lifecycle Management, and Monitoring, Visibility & Control Integration. Each capability was rated by implementation level to pinpoint where controls were consistent and where gaps remained.

Third, we completed an inventory review of active initiatives, documenting backlog items and roadmap artifacts to validate current execution and avoid duplicative recommendations. We also benchmarked the client's identity control set against a consolidated baseline derived from industry best practices and NIST 800-53r5 to assess coverage, consistency, and relative maturity.

Because agentic AI was the driving threat prompting this effort, we extended the analysis beyond traditional NHIs, mapping where LLM, agent, and local execution (MCP) risks converge, defining AI identity best practices across six domains, and evaluating control options spanning corporate, personal, and shadow AI usage.

The result was a prioritized 2026 to 2027 roadmap of 20+ initiatives, sequenced quarter by quarter across the five NHI management domains and validated with client leadership.

Empirically Grounded Findings

10+ SME discovery sessions across the cyber organization ensured recommendations reflected operational reality, not just framework theory.

Structured Maturity Baseline

An assessment at the capability level across five NHI domains gave leadership a clear, defensible view of implementation levels and gaps.

Benchmarking Against Best Practices

Comparison at the control level against industry best practices and NIST 800-53r5 validated coverage and identified where standards needed greater depth.

Threat Coverage Ready for AI

Analysis extended to agentic AI, LLM, and MCP risks: the fastest growing and least governed edge of the identity landscape.

Actionable, Sequenced Roadmap

A quarterly roadmap spanning two years translated findings into 20+ prioritized initiatives aligned to ownership and work.

Outcome

We delivered a defensible NHI operating model and a two year roadmap that positions the bank to govern identities at AI scale

We provided the client with a complete NHI management operating model and control landscape, a maturity baseline backed by evidence, and a prioritized roadmap that leadership can defend in executive, audit, and regulatory conversations.

The engagement gave the organization a shared, structured language for NHI risk, clarifying which capabilities are strong, where identity types fall outside traditional governance boundaries, and how emerging AI identities differ from conventional service accounts in lifecycle, security, and governance needs.

The roadmap moves the program from reactive to deliberate. Initiatives are sequenced across six quarters and five management domains, from governance and inventory expansion through credential rotation automation, lifecycle integration, and unified monitoring. Each is traceable to an identified gap and validated against work already underway, ensuring no duplication of effort.

Critically, the analysis positions the bank ahead of the agentic AI curve. By defining AI identity best practices, mapping converged LLM, agent, and MCP risks, and evaluating control options across managed, unmanaged, and shadow usage, the client now has a framework for extending identity governance to AI systems as adoption scales.

100s

Local LLMs evaluated for IAM risk

20+

Prioritized Initiatives on a Two-Year Roadmap

30k+

NHIs inventoried and assessed

Making the Case for Automation

Manual identity governance can't keep pace with agents that provision access, chain permissions, and act at machine speed. Automation is what makes behavior,-scope and audit-level oversight achievable at scale turning governance from a periodic review into a continuous, real-time control.

WASHINGTON D.C. (Headquarters)

8609 Westwood Center Dr., Suite 110,
Tysons Corner, VA 22182

CONTACT US

Andrew Paek, Consultant
andrew.paek@dayblinkconsulting.com

Michael Morgenstern, Partner
michael.morgenstern@dayblinkconsulting.com

DAYBLINK CONSULTING

www.dayblinkconsulting.com