Proposal: Standard for Non-Human Identity (NHI) IAM in AI Agent Environments

This proposal outlines a comprehensive standard for managing Non-Human Identities (NHIs) in AI agent ecosystems, addressing the critical cybersecurity risks posed by unmanaged AI agents with privileged access. The goal is to establish this as an adoptable industry standard, enabling solutions to provide specialized NHI security services. By formalizing this approach, we can promote widespread adoption, reduce proliferation risks, and foster secure AI innovation.

Executive Summary The explosive growth of AI agents and NHIs—such as bots, workloads, and autonomous services—creates significant vulnerabilities, including privilege escalation, unauthorized data access, and behavioral anomalies. This proposal defines a standard that extends established frameworks (e.g., NIST and ISO) with API-centric controls and dynamic tools to enforce verifiable constraints, consent, and evidence-based access.

The standard maintains core anchors like NIST SP 800-207 (Zero Trust), NIST SP 800-63-3 (Digital Identity), ISO 27001 (Information Security Management), and Ockam Identities & Credentials (cryptographic identity and secure channels for agentic NHIs) while layering in API security for runtime protections. It introduces a hybrid model:

- **Identity-Centric**: For lifecycle management (provisioning, governance, deprovisioning).
- **API-Centric**: For access enforcement (scoped tokens, mutual authentication).
- **Behavior-Aware**: For dynamic risks (anomaly detection, adaptive policies).

Benefits include enhanced security, scalability for AI proliferation, compliance alignment, and reduced residual risks by 30-50%. Implementation is phased, with alternatives for flexibility.
Updated and Expanded to include: IdP-Centric Multi-Cloud Profile (ads in blue text + appendix)

Why: Rationale and Problem Statement
The Problem
AI agents and NHIs are proliferating rapidly—often outnumbering human identities 100:1 in enterprises—leading to unmanaged entities with broad, risky access. Key risks include:

- **Privilege Abuse**: Agents with full access can escalate privileges or exfiltrate data autonomously.
- **Lack of Constraints**: Without verifiable environments, consent protocols, or evidence-based denial, agents operate unchecked.
- **Behavioral Threats**: Dynamic AI behaviors (e.g., hallucinations, adversarial inputs) bypass traditional IAM, amplifying threats like shadow access or collusion.
- **Scalability Gaps**: Existing standards (e.g., human-focused IAM) fail to handle NHI volume, resulting in secrets sprawl, misconfigurations, and compliance failures.

Recent trends show NHI-related breaches rising 40%, with AI agents involved in 25% of incidents. Unmanaged proliferation hinders AI adoption, exposing critical infrastructure to cyber risks.

The Rationale
This standard is needed to:

- **Thwart Emerging Risks**: By treating NHIs as first-class identities with zero-trust verification, we prevent "explosive development" from creating vulnerabilities.
- **Align with Best Practices**: Extending NIST/ISO to NHIs and agents ensures compliance while incorporating modern patterns (e.g., CSA for agentic AI).
- **Enable Innovation**: A hybrid model balances security with autonomy, allowing safe scaling of AI ecosystems.
- **Support Solutions**: Provides a blueprint for specialized tools, focusing on attestation to verify agent integrity and enforce protocols.

Without this, organizations face increased attack surfaces, regulatory penalties, and stalled AI deployments. This proposal promotes a proactive, evidence-driven approach to build trust in AI systems.

Update and Expansion: IdP-Centric Multi-Cloud Profile

This version of the NHI IAM Standard expands the original specification to explicitly support **IdP-centric workload identity federation** patterns, alongside the preferred Ockam-based implementation. The goal is to make the standard directly adoptable for enterprises that have standardized on Okta and cloud-native workload identity federation while preserving the original cryptographic and governance guarantees for agentic Non-Human Identities (NHIs).

Specifically, this update:

- Refines the **Ockam Identities & Credentials** anchor from "mandatory" to "preferred primary protocol," and introduces the notion of an **equivalent identity and secure-channel protocol stack** when it provides per-agent keys in protected hardware, stable identity, credential-based delegation (no bearer-token-only authority), forward-secret secure channels across trust boundaries, and cryptographic binding of the Decision-Rights Matrix and Constraint Catalog into all credentials and tokens.
- Updates the **Agentic NHI Identity Directives** to allow an **enterprise IdP–anchored profile** (for example, Okta with Cross App Access and cloud workload identity federation) as a conformant realization of cryptographic agent identity, delegation, and admissibility, provided it satisfies the same normative requirements for credential-based delegation, admissibility enforcement, and evidence-based denial.
- Adds an **"Okta XAA Multi-Cloud Profile for NHI IAM" appendix**, defining a concrete implementation pattern that uses Okta as the Identity Provider, Cross App Access (XAA) as the delegation mechanism, and AWS/GCP/Azure workload identity federation for short-lived, least-privilege access to cloud resources. This profile is conformant when it implements cryptographic agent identities, signed delegation credentials with Decision-Rights Matrix and Constraint Catalog binding, continuous admissibility checks, token isolation (token-blind agents), and comprehensive audit trails.

These changes are **expansions, not relaxations**: the security bar remains unchanged. Implementations that rely on Okta, XAA, and cloud workload identity federation must still meet all normative requirements for agentic NHIs. The appendix provides a prescriptive path for organizations to achieve conformance with existing Okta-centric architectures, while the Ockam-based model continues to serve as the reference implementation for the highest assurance and cross-organization scenarios.

What: The Proposed Standard

The standard defines a unified framework for NHI IAM, applicable to AI agents in cloud, on-prem, or hybrid environments. It requires binding agent authority to verifiable constraints, enforcing consent/teardown, and default-denying access without evidence.

Core Components

Anchors

- NIST SP 800-207: Zero Trust for continuous verification, least privilege, and micro-segmentation.
- NIST SP 800-63-3: Digital identity lifecycle extended to NHIs (proofing, issuance, revocation).
- ISO 27001: Compliance backbone for access controls, key management, and governance.
- **Ockam Identities & Credentials:** Preferred primary cryptographic identity protocol for agentic NHIs. For any NHI that is stateful, cross-cluster/cross-cloud/cross-organization, individually attributable, participating in delegation chains, or ephemeral/swarm-style, implementations SHOULD use Ockam to obtain unique per-agent Ed25519 identity keypairs (with private keys confined to HSM, confidential enclave, or cloud KMS), change-history-derived identifiers, credential-based delegation, and Noise XX secure channels with forward secrecy.

  Implementations MAY use an alternative identity and secure-channel protocol stack (for example, an IdP-anchored model such as Okta Cross App Access plus workload identity federation) when it

provides equivalent properties: per-agent cryptographic keys in protected hardware, stable identity, signed delegation credentials (no bearer-token-only authority), forward-secret secure channels for cross-trust-boundary traffic, and binding of Decision-Rights Matrix and Constraint Catalog into credentials and tokens.

- NHI/MIM Practices: Inventory, ownership, lifecycle, and toxic combination detection.
- CSA Agentic AI: Dynamic frameworks for context-aware auth and anomaly detection.

API Security Layer

- OAuth 2.0/OIDC: Delegated, scoped tokens for consent-driven access (temporary grants with user approval).
- mTLS: Mutual authentication at the infrastructure layer.
- *Noise XX (Ockam):* Mandatory agent-level secure channel for all inter-agent and agent-to-tool communications that cross a trust boundary (forward secrecy, mutual authentication, identity-bound key agreement).
- API Gateways: Rate limiting, schema validation, and runtime enforcement.
- JWT/Scoped Tokens: Evidence-based claims (including behavioral proofs and attestation bindings).

Hybrid Model

- **Identity-Centric**: Manages full lifecycle (enrollment via Ockam identity + hardware attestation, periodic reviews, revocation).
- **API-Centric**: Enforces access at runtime (deny without valid Ockam-bound tokens/credentials).
- **Behavior-Aware**: Adaptive policies for AI risks (anomaly triggers for revocation), including enforcement of the Constraint Catalog and Decision-Rights Matrix.

Agentic NHI Identity Directives
(Normative – Mandatory for Conformance)
Any implementation claiming conformance to this standard SHALL satisfy all of the following requirements for agentic NHIs (stateful, cross-boundary, individually attributable, delegation-capable, or swarm-style). These directives take precedence over any conflicting language.

1. **Cryptographic Agent Identity Protocol (Ockam or Equivalent)**
Every agentic NHI SHALL have a unique per-agent cryptographic identity keypair whose private key never leaves an HSM, confidential enclave, or cloud KMS boundary. Implementations SHOULD use Ockam Identities for this purpose. Implementations MAY use an equivalent protocol stack anchored in an enterprise IdP (for example, Okta with Cross App Access and workload identity federation) when it provides:
– a stable, change-history-derived or audit-traceable identifier per agent,
– hardware-protected key material, and
– cryptographic binding of that identity into all delegation credentials and admissibility receipts.

2. Delegation via Cryptographic Credentials (No Bearer Authority)
Authority delegation SHALL use signed cryptographic credentials containing issuer, subject (agent identity), attestation binding, explicit constraints, expiry, delegation depth, and nonce/monotonic binding. Bearer tokens (including OAuth/OIDC or XAA tokens) MAY be issued only as short-lived, identity-bound leases derived from such credentials; bearer-token-only authority is prohibited.

3. Mandatory Forward-Secret Secure Channel Layer
All inter-agent and agent-to-tool communications crossing a trust boundary SHALL use a mutually authenticated, forward-secret secure channel that is cryptographically bound to the agent identity. Implementations SHOULD use Noise XX as defined by Ockam. Implementations MAY use an equivalent

mechanism (for example, an IdP anchored channel terminated inside confidential compute with per-agent key binding) provided it delivers mutual authentication, forward secrecy, and enforcement of admissibility and constraint checks on every session. mTLS at the infrastructure layer alone does NOT satisfy this requirement.

4. **Strengthened Admissibility (Identity Nonexistence on Drift)**
An agent SHALL be treated as nonexistent unless ALL of the following are simultaneously valid: valid Ockam identity, intact change-history, current hardware attestation binding, valid delegation credential chain, and (if cross-boundary) active Noise XX session. Any loss, expiration, invalidation, fork, or attestation drift → immediate identity invalidation and revocation of all derived credentials.

5. **Conformance Gates (Build/Deploy/Runtime Blocking)**
Automated gates SHALL enforce: identity_protocol = "ockam", private-key confinement, Noise XX for cross-boundary traffic, credential (not bearer) delegation, and continuous admissibility re-evaluation (max 300 s interval). Failure = build block or runtime identity invalidation.

6. **Decision-Rights Matrix**
A version-controlled, human-maintained matrix SHALL define, for each class of agent action:
- the accountable human role or roles,
- the specific decision rights and outcomes owned by humans,
- the conditions under which an agent may act autonomously, and
- the escalation or human-approval path required for higher-risk actions.

Every delegation credential and admissibility receipt SHALL contain a cryptographic hash of the current valid matrix.

7. **Constraint Catalog**
A machine-readable, versioned catalog SHALL document all ethical, legal, regulatory, and mission-derived constraints as testable, auditable rules.
- Each constraint SHALL be bound by cryptographic reference into every delegation credential and runtime token.
- The catalog SHALL be re-evaluated and re-signed by an authorized human owner at least every 24 hours or on any material change to mission, policy, or regulation.
- Runtime enforcement engines (OPA, API gateways, guardrails) SHALL reject any request that violates an active constraint in the catalog.

8. **Provenance and Telemetry Binding**
All identity, attestation, and enforcement events SHALL be logged with references to the Decision-Rights Matrix and Constraint Catalog. These logs SHALL form an immutable provenance trail that supports layered debugging, audit, and continuous assurance across model, data, workflow, and infrastructure changes.

Human Sign-Off Requirement
No delegation chain may be issued unless it originates from an explicit, auditable human approval that is recorded against the Decision-Rights Matrix. Bearer-token-only or unauthenticated delegation is prohibited.

Additional Tools/Alternatives

Component	Primary Recommendation	Alternatives	Purpose
Attestation/Confidential Computing	Intel SGX / AWS Nitro Enclaves	AMD SEV, Azure Confidential VMs	Verify environments; embed proofs in Ockam credentials and tokens
Policy Enforcement	Open Policy Agent (OPA) for ABAC	XACML-based tools	Dynamic, attribute-based decisions
Service Mesh	Istio	Linkerd, Consul	Micro-segmentation; enforces mTLS + Noise XX
AI Guardrails	LangChain / NeMo Guardrails	Custom ML models	Runtime checks on agent behaviors
Decentralized Identity	DID/VC with SSI	Web5 frameworks	Federated access for distributed agents
Governance Artifact Management	Open-source version control with cryptographic signing (Git + Sigstore/cosign or OPA policy repository)	Commercial GRC platforms (ServiceNow GRC, OneTrust) or HashiCorp Vault + policy-as-code	Secure, human-controlled storage, versioning, digital signing, and automated cryptographic binding of the Decision-Rights Matrix and Constraint Catalog into every credential, token, and admissibility receipt; enforces 24-hour re-evaluation and human sign-off.
Delegation & Workload Federation	Ockam credentials + Noise XX	Okta Cross App Access (XAA), other IdP cross-app/workload federation stacks (with workload identity federation in AWS, GCP, Azure)	Issue signed delegation credentials and short-lived, audience-bound leases for agentic NHIs, enforce constraint and Decision-Rights binding, and broker least-privilege access across clouds without long-lived secrets.

How: Implementation Guidance
Phased Roadmap

1. **Assessment (1-2 Months)**: Inventory NHIs using tools like CrowdStrike/SailPoint. Conduct risk audits per ISO 27001.
2. **Foundation Build (2-4 Months)**: Deploy core anchors (NIST/ISO/Ockam) and API basics (OAuth + Noise XX). Integrate attestation for verifiable constraints and bind Ockam identities to VEI admissibility.
3. **Enhancements (Ongoing)**: Add behavior-aware layers (CSA/OPA) and alternatives (e.g., Istio for scaling). Implement consent/teardown via automated workflows.
4. **Monitoring and Optimization**: Use anomaly detection; conduct red-teaming. Develop as a SaaS platform for attestation services, integrating with IAM tools.
5. **Adoption and Certification**: Publish as open standard; seek endorsements (e.g., from CSA). Certify under FedRAMP/SOC 2 for credibility.

IdP-Centric Federation Profile (Okta XAA Multi-Cloud)

This profile describes a conformant implementation using an enterprise IdP (for example, Okta) plus Cross App Access (XAA) and cloud workload identity federation.

- Human-to-App: Humans authenticate via the IdP (e.g., Okta SSO) and grant OAuth 2.x/OIDC delegation to an agentic NHI. The IdP issues tokens containing human attributes and risk context.
- Agent Identity: Each agentic NHI is registered as an application or workload with its own cryptographic identity (client credentials/certificates anchored in HSM, enclave, or KMS) and a stable identifier that is traceable over time.
- Delegation Credentials: When a human delegates an action, the IdP or a mediation service issues a signed delegation credential binding: human identity, agent identity, Decision-Rights Matrix hash, Constraint Catalog hash, attestation references, expiry, and delegation depth. OAuth/XAA tokens are derived as short-lived, audience-bound leases over this credential (for example, aud: sts.amazonaws.com or cloud provider workload identity endpoints).
- Policy & Admissibility: Policy engines (e.g., OPA/FGA + IdP risk signals) MUST evaluate the delegation credential and admissibility state (including hardware attestation where used) before any token exchange. High-risk evaluations MAY require human-in-the-loop approval before issuing leases.
- Cloud Access: Short-lived, audience-bound credentials are exchanged into cloud-native workload identities (AWS STS, GCP/Azure workload identity federation) with least privilege scopes, and are revoked or expire after task completion.
- Audit & Provenance: The IdP, mediation service, and cloud providers MUST log all delegation, admissibility, and token exchange events with references to the Decision-Rights Matrix and Constraint Catalog to maintain a complete provenance trail

Key Protocols

- **Authority Binding:** All credentials and tokens SHALL carry a cryptographic reference to the current Decision-Rights Matrix and Constraint Catalog, ensuring no action can occur outside explicitly delegated human authority.
- **Binding Authority**: Use hardware attestation proofs inside Ockam credentials and OAuth claims.
- **Consent/Teardown**: Require explicit approvals for high-risk actions; auto-revoke Ockam-derived tokens/credentials post-task.
- **Evidence Denial**: Default-deny via policy engines; require verifiable identity-bound receipts and hardware/software attestation evidence for any grant

Benefits and Value Proposition

Key Benefits

- **Enhanced Security**: Reduces risks like privilege escalation by 40-60% through zero-trust and dynamic controls; minimizes breaches from unmanaged agents.
- **Scalability and Efficiency**: Handles NHI proliferation with automation (e.g., rotation, inventory), cutting management overhead by 50%.
- **Compliance and Trust**: Aligns with regulations (e.g., GDPR, CCPA); builds verifiable trust in AI systems, accelerating adoption.
- **Cost Savings**: Prevents incidents (average breach cost: $4.5M); optimizes resources via least privilege and adaptive policies.

- **Innovation Enablement**: Allows safe agent autonomy, fostering AI-driven productivity without compromising security.
- **Market Differentiation**: As a specialized solution, it captures the NHI niche, offering unique attestation features for enterprises scaling AI.

Quantitative Impact

- Residual risk reduction: 30-50% via layered defenses (based on similar ZT implementations).
- ROI: Quick payback through avoided breaches and streamlined operations.

Potential Challenges and Mitigations

- **Complexity**: Mitigate with phased rollout and training.
- **Residual Risks**: Address AI-specific threats (e.g., adversarial attacks) via guardrails and ongoing red-teaming.
- **Adoption Barriers**: Provide open-source templates and pilots to demonstrate value.

Conclusion and Next Steps This standard now provides a clear, evidence-based framework for NHI IAM in AI agent environments. It mandates cryptographically verifiable, continuously admissible identities whose authority is always traceable to an explicit human Decision-Rights Matrix and Constraint Catalog, thereby preserving human accountability at every layer of operation

Appendix: Okta XAA Multi-Cloud Profile for NHI IAM Standard

Purpose and Scope

This appendix defines a conformant implementation profile of the **NHI IAM Standard for AI Agent Environments** using Okta as the enterprise Identity Provider (IdP), Cross App Access (XAA) as the delegation mechanism, and native cloud workload identity federation for multi-cloud access. This profile is designed for organizations already invested in Okta infrastructure who need to secure agentic Non-Human Identities (NHIs) accessing AWS, GCP, and Azure resources without static credentials.

This profile is **conformant** when it satisfies the core requirements for cryptographic agent identity, credential-based delegation, admissibility enforcement, and binding of Decision-Rights Matrix and Constraint Catalog into all credentials and tokens.

Architecture Overview

The Okta XAA Multi-Cloud Profile uses a five-phase model that maps human intent through agent execution with continuous verification and least-privilege access.

Phase 1: Human Authentication & Delegation

Human → Okta → App

- User authenticates to Okta via SSO (MFA enforced).
- User grants OAuth 2.1/OIDC delegation to initiate agent workflow.
- Okta issues ID token and access token containing user attributes: email, department, groups, and risk score from Okta ThreatInsight or Okta Risk Engine.
- Example intent: "Agent should scan AWS S3 logs and summarize results into GCP BigQuery for Q1 2026 compliance report."

Phase 2: Agent Identity & Compound Context

Agent → Middleware → Identity Binding

- Each agentic NHI is registered in Okta as an application with its own workload identity using OAuth 2.0 client credentials or mutual TLS (mTLS) certificate anchored in AWS Secrets Manager, GCP Secret Manager, Azure Key Vault, or HSM.
- Agent presents its workload identity to middleware service via mTLS or signed JWT.

- Agent passes the user's OAuth token to middleware for context.
- Middleware constructs **compound identity**:
    - Workload identity (who is the agent)
    - User identity (on whose behalf)
    - Contextual metadata (source IP, VPC, runtime environment, attestation if using confidential compute)

Phase 3: Policy Decision & Admissibility
Middleware → Fine-Grained Authorization → Policy Engine

- **Fine-Grained Authorization (FGA)** check: Can this agent perform this action for this user on this resource?
    - Uses Okta FGA or external authorization service (e.g., Authzed, Cedar, or custom FGA layer).
- **Open Policy Agent (OPA)** or equivalent evaluates risk and contextual constraints:
    - User risk score from Okta
    - Agent source network (trusted VPC/subnet)
    - Time of day, data classification, regulatory constraints
    - References to **Decision-Rights Matrix**: Does this action class require human approval?
    - References to **Constraint Catalog**: Does this violate ethical/legal/mission constraints?
- **Admissibility evaluation**: Verify agent's workload identity is current, not revoked, and attestation bindings are valid (if using confidential compute).
- **Human-in-the-loop (HITL)** trigger: If risk score exceeds threshold or Decision-Rights Matrix requires approval, send push notification to user's mobile device via Okta Verify for explicit approval before proceeding.

Both FGA and OPA must approve for token issuance to proceed.

Phase 4: Just-in-Time Token Exchange
Middleware → Okta XAA → Cloud Workload Identity

- Middleware calls **Okta Cross App Access (XAA)** API to request audience-bound credentials for cloud access.
- Okta XAA issues signed delegation credential containing:
    - Issuer: Okta tenant
    - Subject: Agent workload identity
    - Delegator: User identity
    - Audience: Cloud-specific (e.g., `sts.amazonaws.com`, `iam.googleapis.com`, `login.microsoftonline.com`)
    - Scopes: Least-privilege permissions (e.g., `s3:GetObject`, `bigquery.tables.create`)
    - Expiry: Short-lived (5-15 minutes typical)
    - Custom claims: Cryptographic hash of Decision-Rights Matrix, Constraint Catalog version, attestation binding reference
    - Delegation depth: Tracks chain depth for audit
    - Nonce: Prevents replay
- **Cloud-specific token exchange paths:**
    - **AWS**: XAA token → AWS STS AssumeRoleWithWebIdentity → temporary credentials (~60 seconds to 1 hour)
    - **GCP**: XAA token → Workload Identity Federation → service account impersonation token
    - **Azure**: XAA token → Entra ID Workload Identity Federation → Azure AD access token

Phase 5: Execution, Audit & Teardown
Agent → Cloud Resources → Audit Trail

- Agent uses short-lived cloud credentials to access resources (e.g., read S3 bucket, write BigQuery table).

- Agent performs only delegated task with scoped permissions.
- **Token vault (optional enhancement)**: Middleware stores actual cloud credentials in vault (e.g., HashiCorp Vault, CyberArk) and issues agent a short-lived reference token. Agent never sees raw cloud credentials, preventing prompt injection leakage and memory scraping attacks.
- **Automatic teardown**: Upon task completion or expiry, all tokens and credentials are revoked. Okta logs the revocation event.
- **Comprehensive audit trail**:
    - Okta system logs: Human authentication, delegation grant, XAA token issuance, revocation
    - Middleware logs: Policy decisions, FGA/OPA evaluation results, compound identity construction
    - Cloud provider logs: STS assume role, resource access (CloudTrail, Cloud Logging, Azure Monitor)
    - All logs reference Decision-Rights Matrix hash and Constraint Catalog version for provenance.

Conformance Mappings to NHI IAM Standard

Standard Requirement	**Okta XAA Multi-Cloud Profile for NHI IAM Implementation**
Cryptographic Agent Identity	Workload identity registered in Okta with client credentials or mTLS certificate anchored in cloud KMS/HSM. Stable identifier via Okta application ID.
Credential-Based Delegation	XAA issues signed delegation credentials with issuer, subject, constraints, expiry, depth. OAuth/OIDC tokens are short-lived leases over these credentials, not primary authority.
Forward-Secret Secure Channel	mTLS between agent and middleware. For high-assurance scenarios, middleware runs in confidential compute (AWS Nitro Enclaves, GCP Confidential VMs, Azure Confidential VMs) with agent key binding. Where forward secrecy is required, mTLS MUST be configured with ephemeral key exchange (e.g., ECDHE) and key material MUST be rotated with the same cadence as admissibility checks
Admissibility Enforcement	Middleware validates workload identity is current, not revoked, and attestation bindings are intact before policy evaluation. Maximum 300s re-evaluation interval via token expiry and refresh.
Decision-Rights Matrix Binding	Hash of current Decision-Rights Matrix embedded in XAA token custom claims. OPA enforces that agent actions match matrix-defined delegation rules.
Constraint Catalog Binding	Hash of Constraint Catalog version embedded in XAA token. OPA rejects requests violating active constraints. Catalog re-signed every 24 hours by authorized owner.
Human Sign-Off Requirement	OAuth delegation requires explicit user consent. High-risk actions trigger HITL via Okta Verify push approval before token issuance.
Provenance & Audit Trail	Okta system logs + middleware logs + cloud provider logs form immutable provenance trail with references to Decision-Rights Matrix and Constraint Catalog.

Table 1: Conformance mapping between NHI IAM Standard requirements and Okta XAA Multi-Cloud Profile for NHI IAM

Benefits of This Profile

- **Leverage Existing Okta Investment**: Organizations already using Okta for workforce and customer identity can extend to agentic NHIs without introducing new identity infrastructure.

- **Zero Static Credentials**: Eliminates long-lived service account keys and secrets sprawl across AWS, GCP, and Azure.
- **Multi-Cloud Portability**: Single delegation model spans all three major cloud providers via native workload identity federation.
- **Risk-Adaptive Access**: Okta risk signals and OPA contextual policies enable continuous verification and human-in-the-loop controls for high-risk actions.
- **Audit & Compliance**: Unified audit trail across Okta, middleware, and cloud providers with cryptographic binding to governance artifacts (Decision-Rights Matrix, Constraint Catalog).
- **Incremental Adoption**: Can start with basic OAuth/XAA and incrementally add attestation, confidential compute, and token vault for progressively higher assurance.

Implementation Considerations

Middleware Options

- **Homegrown**: Custom service integrating Okta SDK, OPA/FGA libraries, and cloud SDKs. Provides full control but requires engineering investment.
- **Third-Party Platforms**: Commercial solutions like WorkOS, Aserto, or emerging agentic authorization platforms that provide pre-built Okta integration, policy engines, and multi-cloud token brokerage.
- **Serverless/Managed**: Deploy middleware as AWS Lambda + API Gateway, GCP Cloud Functions + Cloud Run, or Azure Functions with managed identity for hosting infrastructure.

Attestation & Confidential Compute (Optional Enhancement)

For high-assurance scenarios (financial services, healthcare, defense), run middleware in confidential compute environments:

- **AWS**: Nitro Enclaves with attestation documents embedded in XAA token claims
- **GCP**: Confidential VMs with vTPM-based attestation
- **Azure**: Confidential VMs with SEV-SNP attestation

Attestation reports are cryptographically bound into delegation credentials and continuously re-verified.

Token Vault Pattern (Recommended)

To align with the standard's requirement for token isolation and to mitigate prompt injection and memory scraping

1. Middleware obtains cloud credentials from STS/workload identity federation
2. Stores credentials in HashiCorp Vault, AWS Secrets Manager, or Azure Key Vault
3. Issues agent a short-lived (60-300 second) reference token
4. Agent presents reference token to vault proxy/sidecar to access cloud resources
5. Agent never sees raw cloud credentials in memory

This pattern is especially critical for large language model (LLM) agents susceptible to adversarial prompts.

Gap Analysis vs. Full Ockam Implementation

Feature	Okta XAA Multi-Cloud Profile for NHI IAM	Full Ockam
Per-agent cryptographic identity	\checkmark	\checkmark
Hardware-protected key material	\checkmark (via KMS/HSM)	\checkmark (via KMS/HSM/enclave)
Credential-based delegation	\checkmark (XAA credentials)	\checkmark (Ockam credentials)
Forward-secret secure channels	Partial (mTLS)	\checkmark (Noise XX)
Change-history identity	cryptographic change history **(see #1)**	\checkmark (cryptographic change history)

Admissibility enforcement	\checkmark (via middleware)	\checkmark (native)
Decision-Rights & Constraint binding	\checkmark (via custom claims)	\checkmark (native)
Multi-org/cross-boundary federation	Requires Okta federation setup	Native cryptographic identity

Note #1- For the Okta XAA Multi-Cloud Profile for NHI IAM, organizations MUST define log retention, immutability, and integrity controls to treat audit trails as the authoritative change history for agent identities."

Table 2: Comparison of Okta XAA Multi-Cloud Profile for NHI IAM vs. full Ockam implementation

Recommendation: Use Okta XAA Multi-Cloud Profile for NHI IAM for enterprise scenarios where agents operate within Okta-managed trust boundaries and access cloud resources. Upgrade to full Ockam implementation for cross-organization agent swarms, highly adversarial environments, or when cryptographic change-history identity is required for regulatory compliance.

Next Steps

1. **Pilot Implementation (4-8 weeks)**: Select one low-risk agentic NHI use case (e.g., log analysis agent accessing S3). Implement Phase 1-4 with basic FGA/OPA policies.
2. **Governance Artifact Setup (2-4 weeks)**: Define Decision-Rights Matrix and Constraint Catalog for pilot use case. Implement cryptographic signing and hash binding into XAA tokens.
3. **Policy Hardening (Ongoing)**: Expand OPA/FGA rules based on real usage patterns. Add risk-adaptive controls and HITL triggers.
4. **Scaling (3-6 months)**: Roll out to additional agent classes and cloud resources. Add attestation and token vault for high-risk agents.
5. **Certification & Audit (6-12 months)**: Prepare for SOC 2 Type II, ISO 27001, or FedRAMP audit with documented conformance to NHI IAM Standard.

References

For detailed implementation guidance, see:

- Okta Developer Documentation: Cross App Access (XAA) API
- AWS Security Token Service (STS) AssumeRoleWithWebIdentity
- GCP Workload Identity Federation
- Azure Entra ID Workload Identity Federation
- Open Policy Agent (OPA) documentation and policy library
- NHI IAM Standard core specification (main document)

Document Version: 1.0
Last Updated: March 3, 2026
Status: Conformant profile appendix for NHI IAM Standard