OASIS

White Paper

Why Identity Governance Alone Is Not Enough For Securely Managing Non-Human Identities

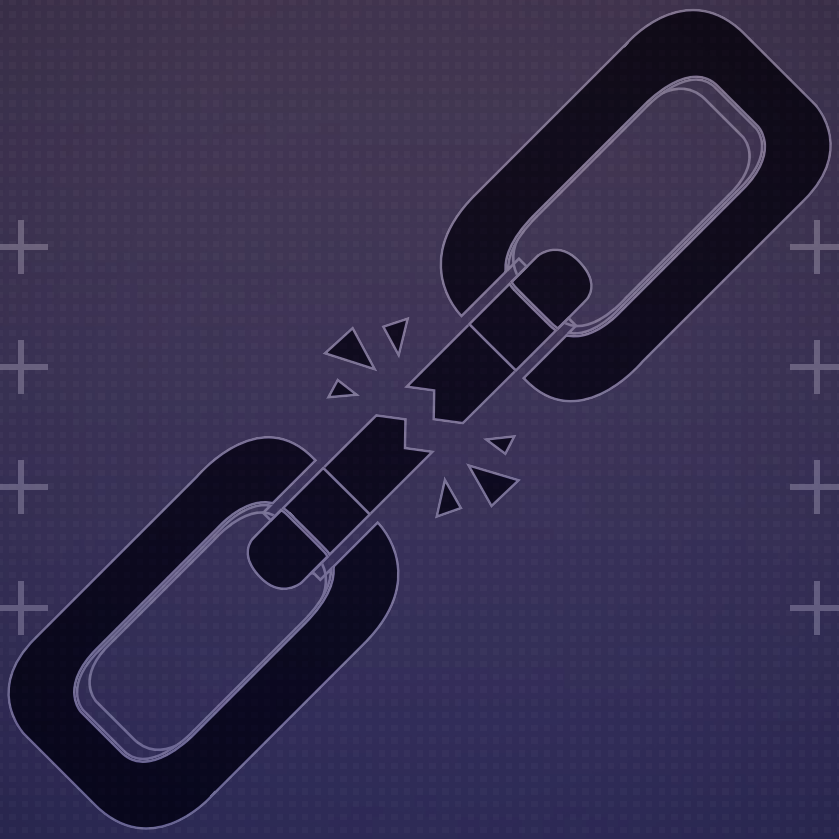

- **Learn** *why traditional IGA platforms fall short for securing non-human identities, and why governing machine and Agentic AI access requires deep context and continuous validation—not static reviews—to ensure both least privilege and business continuity.*
- **Learn** *how the lifecycle between human and machine identities differ and why solutions architectured for machine identities enable more secure provisioning, continuous governance, and confident decommissioning at scale.*
- **Learn** *how to achieve faster time-to-value for your governance program with broader integrations and less effort.*

The Growing Gap Between Identity Governance and Reality

Identity Governance and Administration (IGA) platforms have long served as the foundation for managing human access to enterprise systems. They excel at governing employees and contractors whose identities are anchored in authoritative sources such as HR systems, follow predictable joiner–mover–leaver lifecycles, and require periodic access certification.

However, the identity landscape has fundamentally changed.

Today, the vast majority of identities operating inside modern enterprises are non-human: service accounts, cloud service principals, automation bots, and increasingly, autonomous AI agents. These identities operate continuously, are created and modified outside centralized workflows, and authenticate using secrets, tokens, and certificates rather than passwords with the additional layer of strong protection from MFA.

Many organizations are being told that their existing IGA platforms will soon "do it all" — that with incremental roadmap enhancements, the same tools designed for human governance can adequately discover, govern, and manage the full lifecycle of non-human identities (NHIs).

The limitations, however, are not a matter of missing features. They are architectural.

Governing non-human identities requires a fundamentally different data model, control plane, and lifecycle approach — one that cannot be retrofitted into human-centric IGA systems without significant tradeoffs in accuracy, safety, and time-to-value.

Why Non-Human Identities Break the IGA Model

Hint: Context is Critical

IGA platforms were designed around a core assumption: identities represent people.

Human identities have:

- A clear authoritative source (HR, contractor systems)
- Stable attributes (department, manager, role)
- Predictable lifecycle events (joiner, mover, leaver)
- Access patterns that change infrequently

Non-human identities violate many of these assumptions as they are often decentralized, have unclear ownership and do not authenticate as human users do. Because NHIs follow different access patterns, context is critical for organizations to evaluate revoking access in a manner that prevents outages from occurring or other major disruptions to business continuity.

NHIs are:

- Created by numerous personas in an organization, platforms, automation pipelines, and third-party integrations
- Rarely tied to a single owner or team
- Sometimes short-lived and subject to change
- Dependent on secrets, tokens, and certificates
- Often shared across systems and workflows
- Created in an ad-hoc manner and may stop being used for numerous reasons

Looking at it in this light, the only thing human and non-human identities share is the word “identity”.

Applying a human governance model to NHIs introduces blind spots that lead to over-provisioning, audit failures, operational outages, and elevated breach risk.

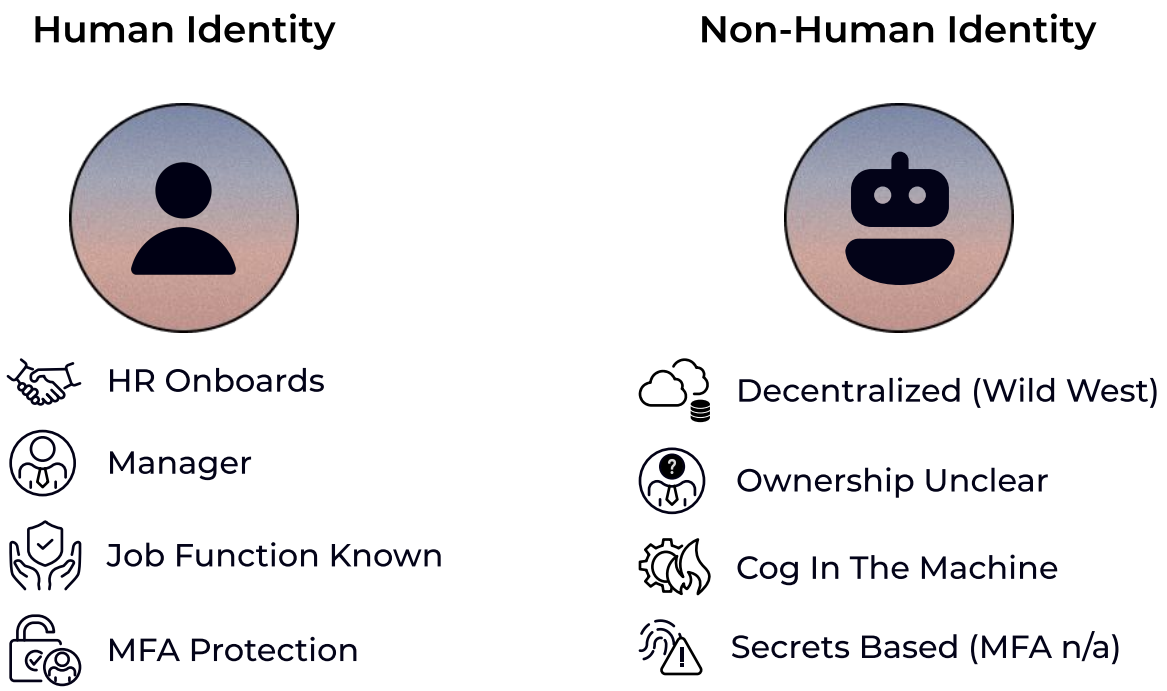

The False Promise of "IGA Can Do It All"

As non-human identities have grown in scale and visibility, many IGA providers have responded by extending their messaging and roadmaps to include NHIs.

In practice, this typically means:

- Aggregating service accounts or cloud principals as "identities"
- Applying static attributes or naming conventions
- Running certification campaigns without usage context

What is often missing is the ability to understand:

- Who or what actually uses an identity
- Which secrets enable that access
- What systems depend on it
- What will break if it changes
- Who owns an identity and its business justification

The issue is not intent — it is design. Lightweight features offered within an IGA suite that focus on NHIs do not eliminate the underlying mismatch between IGA architectures and non-human identity realities.

The Missing Data Model: Why Context Matters

Traditional Identity Governance and Administration (IGA) platforms model access through a simple relationship:

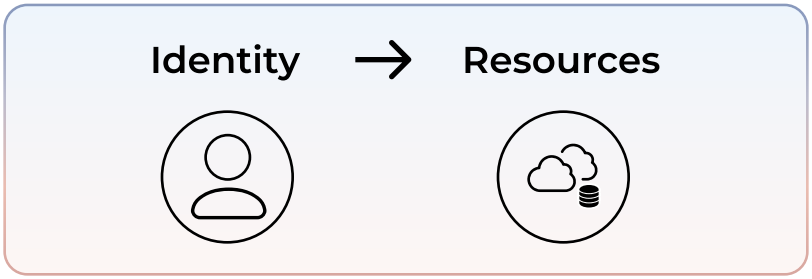

This abstraction works for human users because much of the context required to govern them is already available through business systems — job role, department, manager, and employment status or other user behavior analytics. These signals provide a shared frame of reference for evaluating access, even when permissions are complex.

Non-human identities also require context, but they lack these inherent organizational signals or baselines for behavior. A service account, API key, or workload identity has no role or manager. It doesn't have a location or working hours. Its legitimacy is defined by how it is used in production, not by how it appears in a directory.

As a result, effective non-human identity governance depends on understanding a broader chain of relationships:

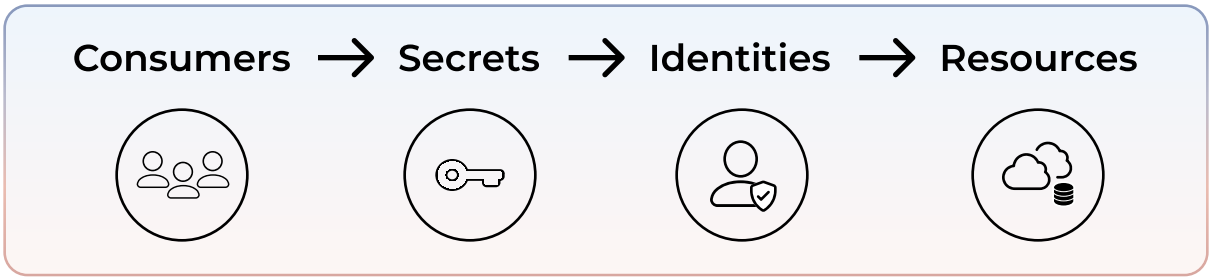

For non-human identities, context does not come from organizational attributes. It comes from how access is established and exercised at runtime.

Every non-human access event follows a common pattern. A consumer, such as an application, service, cloud workload, automation pipeline, or AI agent, needs to perform an action. To do so, it presents a secret, such as a key, token, or certificate. That secret authenticates the consumer as a specific identity, such as a service account, role, or cloud principal. That identity is then authorized to access one or more resources, including APIs, databases, SaaS applications, or cloud services.

Each of these elements is familiar on its own, but none of them are meaningful in isolation. A service account without knowledge of what uses it is indistinguishable from an orphaned identity. A secret without visibility into where it is deployed cannot be rotated safely. A permission without insight into how it is exercised cannot be reviewed with confidence.

For non-human identities, this chain of relationships is the context. Governing NHIs requires visibility into how consumers, secrets, identities, and resources connect in practice, not just how they are defined in a directory.

Traditional IGA platforms struggle here because they were not designed to model or maintain these relationships. They typically surface identities and entitlements in isolation, requiring manual ownership assignments and periodic reviews that offer little insight into actual usage. Without visibility into consumers and secrets, teams are forced to choose between governance rigor and operational safety—often approving access rather than risking disruption.

This lack of context has practical consequences. Organizations cannot reliably distinguish active identities from dormant ones, rotate credentials with confidence, or decommission accounts without fear of breaking production systems. Access reviews become exercises in validation rather than risk reduction.

Reading identities from a directory is straightforward. Understanding why they exist and how they function in production is not.

IGA systems were built to answer: “Who has access to what?” Non-human identity governance must answer: “What is acting, how is it authenticated, and what function does this access serve?”

Without a data model that captures consumers, secrets, identities, and resources as a connected system, governance of non-human identities will remain incomplete—regardless of how mature an organization’s IGA program may be.

Visibility Gap

IGA assumes a single human user for each account. For example, multiple AD service accounts may be synced to a single Entra ID. An IGA platform will not show this relationship and the appropriate context linking these accounts because the model doesn’t support the existence of two ‘identities’ for one account.

This is where purpose-built, non-human identity governance can provide value and visibility. With Oasis, organizations gain deeper understanding into the context of accounts.

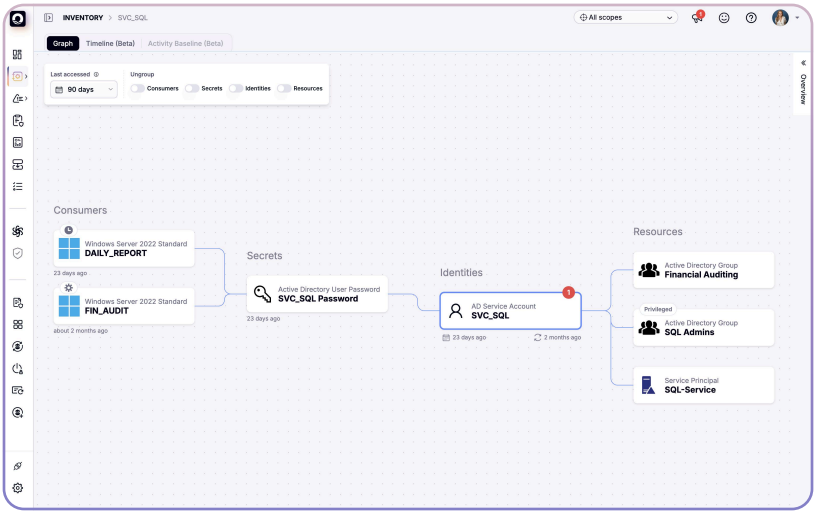

Oasis provides deep context for the relationships and dependencies between consumers, secrets, identities and resources.

Supporting a Fundamentally Different Identity Lifecycle

Discovery and governance are necessary — but on their own, they are insufficient.

Most security programs encounter non-human identities after they already exist: once service accounts have been created, secrets distributed, and access embedded into production workflows. At that point, organizations are forced into a reactive posture — discovering identities retroactively and attempting to govern them without breaking dependent systems.

This is the opposite of how risk should be managed.

Non-human identities require proactive lifecycle management, beginning at creation and continuing through every phase of their existence. Unlike human identities, which can be centrally provisioned based on authoritative sources, non-human identities are created in highly decentralized ways: by developers, cloud platforms, automation tools, CI/CD pipelines, and third-party integrations. They are often instantiated with excessive privileges, shared credentials, and no durable ownership model.

When these identities are created without guardrails, discovery becomes damage assessment rather than governance.

A purpose-built non-human identity lifecycle introduces controls upstream, not just downstream. This includes the ability to:

- Provision non-human identities in a standardized, policy-driven way
- Enforce least privilege at creation, not months later
- Bind identities to explicit owners, use cases, and expiration expectations
- Generate and manage secrets as part of identity creation rather than as a separate, manual step

By establishing these controls at provisioning time, organizations dramatically reduce the number of orphaned, overprivileged, and unknown identities entering their environments in the first place.

Lifecycle management must then continue beyond creation. Non-human identities must be continuously evaluated based on usage and behavior, not static attributes. As applications evolve, integrations change, and automation expands, access requirements shift accordingly. Effective lifecycle management detects these changes and adjusts access, rotates credentials, and decommissions identities safely when they are no longer needed.

This continuous approach stands in contrast to traditional identity governance models that rely on periodic discovery and certification. For non-human identities, waiting for a quarterly review cycle is not just inefficient — it is dangerous.

Proactive provisioning combined with continuous lifecycle enforcement transforms non-human identity governance from a reactive cleanup exercise into a preventative security control. It reduces risk, accelerates time to value, and enables organizations to scale automation and AI initiatives with confidence.

The machine-identity lifecycle may include:

- Creation
- Usage
- Policy enforcement
- Change
- Rotation
- Decommission

Human-centric lifecycle models assume centralized provisioning and infrequent change. NHIs are created dynamically, often at scale, and frequently modified by automation.

Attempting to force NHIs into ticket-based or approval-heavy provisioning workflows introduces friction that slows development and encourages workarounds. Over time, this leads to unmanaged identities outside governance entirely.

True NHI lifecycle management requires:

- Usage-aware provisioning
- Continuous monitoring
- Policy-driven remediation
- Safe, dependency-aware decommissioning

These controls must operate continuously — not only quarterly during certification campaigns.

Rethinking Attestation: We Can Do Better Than Access Reviews

Traditional access certifications were designed to answer a single, human-centric question:
Does this person still need this access?

That question assumes stable roles, clear ownership, and relatively static usage patterns. For non-human identities, those assumptions no longer hold.

Attesting non-human identities requires answering a very different set of questions:

- Is this identity actively in use?
- What workloads, services, or processes depend on it?
- What data and systems does it access?
- What would break if it were changed or removed?
- Who is accountable for its existence and behavior?
- What risk would it pose if compromised?
- What access does it actually need to perform its function?

Traditional IGA platforms are poorly suited to answer these questions. They rely on periodic, manual reviews of identity-to-resource mappings without sufficient visibility into consumers, secrets, or runtime usage. As a result, reviewers are asked to certify access they cannot meaningfully evaluate. The predictable outcome is rubber-stamping, reviewer fatigue, and increasing skepticism from auditors—not because reviews are missing, but because they lack substance.

Human access reviews are typically conducted on a fixed schedule—monthly, quarterly, or annually. While this cadence may satisfy audit requirements, it does not reflect how access risk actually emerges or how least-privilege should be enforced. For non-human identities, where access changes continuously through deployments and automation, scheduled reviews are not always the best approach.

Scale makes this problem significantly worse. In most environments, non-human identities outnumber human users by orders of magnitude. Applying human-oriented attestation processes to machine identities is operationally unsustainable and does little to reduce risk.

A more effective approach treats attestation as a continuous, evidence-driven process rather than a periodic approval exercise. After a non-human identity is provisioned, governance shifts from manual review to policy-based enforcement informed by real usage and dependency context. Examples include:

- Detecting and remediating over-privileged identities
- Identifying credentials that are nearing expiration or no longer used
- Rotating secrets automatically when employees are offboarded or risk conditions change
- Flagging identities that are over-consumed, misused, or behaving anomalously
- Decommissioning identities when their consumers are retired or become inactive

In this model, accountability is established through observable behavior and ownership signals derived from how identities are used in production—not through forced, periodic attestations that provide little assurance.

For non-human identities, effective attestation is not about certifying access on a schedule. It is about continuously validating that access remains necessary, appropriate, and safe, based on evidence. Anything less gives the appearance of governance without delivering real control.

Where Purpose-Built Non-Human Identity Governance Delivers

Non-human identities require a dedicated governance layer — one designed from the ground up for scale, context, and continuous lifecycle management.

Purpose-built NHI governance platforms provide:

- **Deep discovery across on-prem, cloud, SaaS, and automation environments**
- **Contextual mapping of consumers, secrets, identities, and resources**
- **Continuous risk assessment and prioritization**
- **Safe provisioning, rotation, and decommissioning**
- **Readiness for agentic and autonomous identities**

Traditional IGA platforms are poorly suited to answer these questions. They rely on periodic, manual reviews of identity-to-resource mappings without sufficient visibility into consumers, secrets, or runtime usage. As a result, reviewers are asked to certify access they cannot meaningfully evaluate. The predictable outcome is rubber-stamping, reviewer fatigue, and increasing skepticism from auditors—not because reviews are missing, but because they lack substance.

Benefit: Faster Time to Value Without Per-app Connectors

A major advantage of managing non-human identities (NHIs) with a specialized architecture is how quickly organizations can realize value—without the overhead of traditional IGA deployments.

Legacy IGA approaches rely on per-application connectors, forcing teams to integrate each app individually and negotiate onboarding with application owners. In practice, this model is slow, costly, and widely viewed by IGA leaders as one of the most painful aspects of governance programs.

Instead of connecting to every application, this architecture integrates with identity providers, logs, and EDR, delivering broad visibility into non-human identity activity with far fewer integrations. While this approach does not enumerate every local account inside every application, it addresses the majority of the risks organizations need to manage—quickly and with significantly less friction.

This is where Oasis fits.

Rather than attempting to extend human-centric models, Oasis was built specifically to govern non-human identities across their entire lifecycle — from creation to decommissioning — with the contextual intelligence required to operate safely at scale.

IGA platforms remain essential for human identity governance. Oasis complements them by addressing the identities they were never designed to manage with complete governance for NHIs from creation to decommission.

Next Steps

The challenge facing organizations today is not whether identity governance is important — it is whether existing tools are aligned with modern identity realities.

Non-human identities already dominate enterprise environments. Agentic AI will accelerate that shift further. Waiting for incremental roadmap promises increases risk while extending operational pain.

Effective identity security now requires recognizing that humans and machines are governed differently — and adopting architectures designed accordingly.

Learn more about the Oasis platform at www.oasis.security.