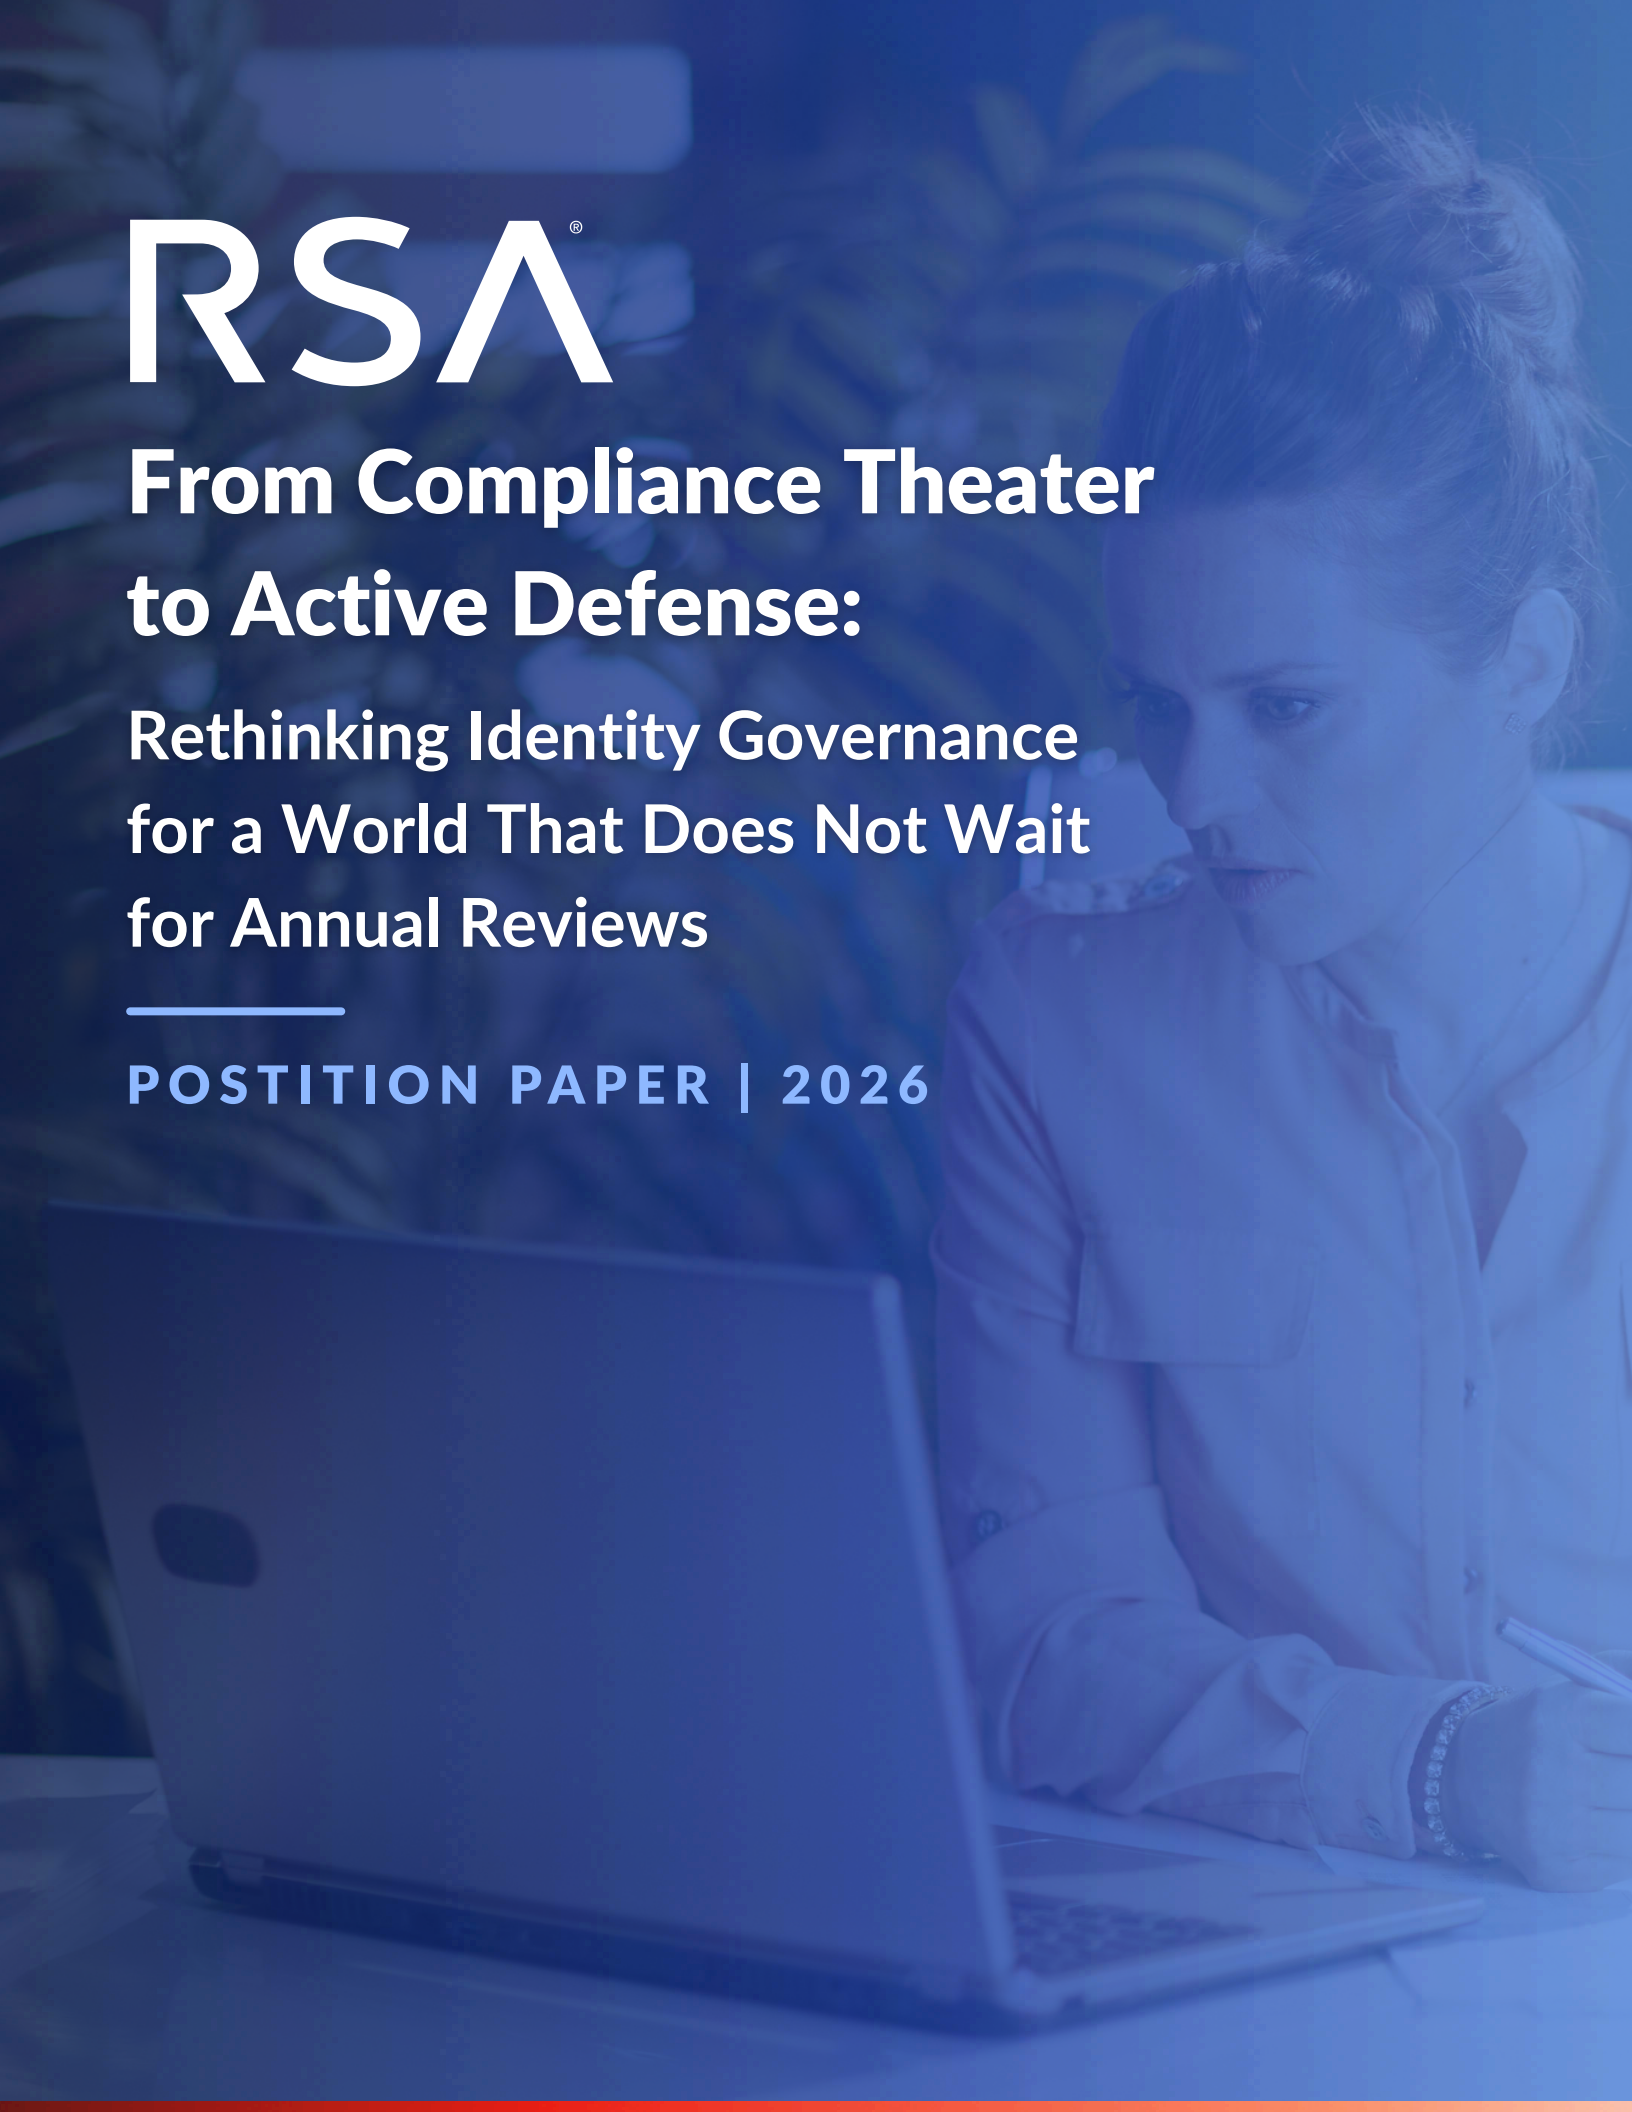

RSA®

From Compliance Theater to Active Defense:

Rethinking Identity Governance for a World That Does Not Wait for Annual Reviews

POSTITION PAPER | 2026

Table of Contents

Executive Summary

Identity governance programs across industries are operating under a flawed assumption: that periodic access reviews, documented policies, and clean audit reports constitute a defensible security posture. They do not.

The threat landscape has fundamentally changed. Identities, human and non-human alike, are provisioned, modified, and exploited in milliseconds. Access rights drift continuously between review cycles. And when a breach occurs, auditors, regulators, examiners, insurers, and litigators do not ask whether policies existed. They ask whether those policies were enforced, in real time, at the moment the incident occurred.

This paper argues that the traditional model of identity governance and administration (IGA), anchored to annual or quarterly review campaigns, creates what we term the 'Negligence Gap': a widening chasm between what an organization's static policy documentation says and what its dynamic systems are actually doing. Left unaddressed, this gap is not merely a security risk. It is a liability, and for financial institutions and government agencies, it is also an audit finding waiting to happen.

The alternative is a shift to Active Defense Governance (ADG), a model in which governance functions as a continuous, automated, risk-based control plane rather than a periodic attestation exercise. When implemented effectively, ADG converts identity governance from a compliance cost center into an operational security capability and, critically, into a source of litigation-ready defensibility.

This position paper intends to provide practitioners, program leaders, and executive stakeholders with a framework for evaluating and evolving their IGA posture beyond periodic compliance activities toward a continuous, defensible ADG model. In the following sections, this paper will detail:

- Why the traditional IGA review cycle is insufficient for today's risks
- How the resulting governance gap creates legal and regulatory exposure
- A framework for implementing ADG, from foundational capabilities to full defensibility
- The metrics organizations can use to measure ADG's effectiveness
- Implementation considerations for organizations at any stage of governance maturity

1. The Compliance Theater Problem

1.1 The clean audit paradox

Organizations today routinely achieve strong results on SOC 2 Type II audits, GDPR readiness assessments, DORA compliance reviews, and CJIS security audits, only to suffer identity-based breaches within weeks of receiving that clean bill of health. This is not a coincidence. It is a structural consequence of how traditional compliance frameworks are designed and how organizations respond to them.

Compliance frameworks, by their nature, assess a given point in time. A SOC 2 engagement evaluates whether controls were in place and operating effectively during a defined period. A CJIS audit validates that access policies governing criminal justice information are documented and followed.

The frameworks differ, but they share the same limitation: a given period, a snapshot of control, not continuous enforcement. Financial institutions under FFIEC guidance, NYDFS Part 500, or PCI DSS v4.0 face auditors increasingly asking for evidence of continuous controls, not point-in-time attestations. Federal agencies subject to FISMA, FedRAMP, or NIST SP 800-53 operate under governance models that document controls at authorization time, but that documentation may not reflect operational enforcement between review cycles. In each case, the audit framework captures a moment. The threat environment does not pause for the review cycle.

The result is what practitioners are increasingly calling the clean audit paradox, or the compliance paradox: the more confidently an organization believes its audit results, the more exposed it may actually be. A green dashboard generated from data that is three, six, or eleven months old is not evidence of control. It is evidence of past control, and in a breach scenario, that distinction is everything.

If your organization experienced a major identity-based breach today, investigators would not ask whether you had a governance policy. They would ask whether that policy was in effect and enforced at the moment the breach occurred. If the answer is “we review that annually,” you are not just insecure. You are indefensible.

> *In the event of a breach, investigators would not ask whether you had a governance policy. They would ask whether that policy was in effect and enforced at the moment the breach occured. If the answer is “we review that annually,” you are not just insecure. You are indefensible.”*

1.2 Audit decay: the hidden expiration date on every review

Every access certification campaign, entitlement review, and role attestation exercise has an expiration date that no one documents. The moment a campaign closes, the data it certified begins to drift from reality. We call this audit decay.

In modern enterprise environments, identities and their associated access rights are not static. Service accounts are created and misconfigured. Contractors are onboarded and their access scopes expand beyond initial intent. Employees change roles and accumulate entitlements that were never revoked. Non-human identities, including API keys, machine accounts, Agentic AI, and Robotic Process Automation (RPA) bots multiply across infrastructure with minimal oversight.

In this environment, a quarterly review cycle means that access drift goes undetected for up to ninety days before it is even eligible to be caught. An annual review means up to a full year of unmonitored exposure. The adversary community is well aware of this dynamic. Dwell times for identity-based attacks are measured in weeks and months, not hours, precisely because traditional governance cadences create predictable windows of opportunity.

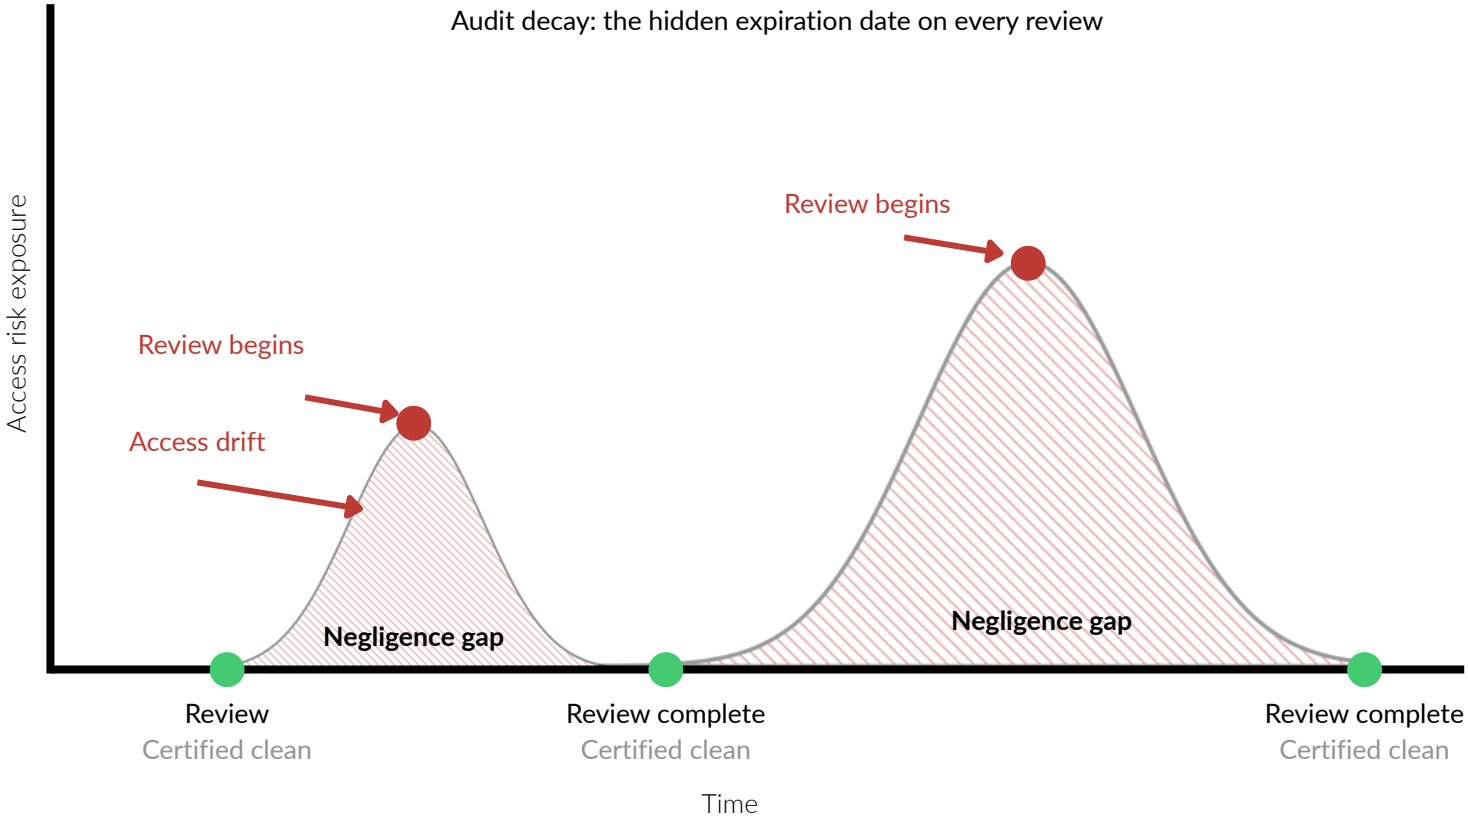

1.3 The supervisor review: governance's most unreliable instrument

The most common form of access review in enterprise IGA programs is the supervisor or manager certification: a periodic request in which managers are asked to review their direct reports' entitlements and confirm that access remains appropriate. This practice is deeply embedded in compliance frameworks and audit expectations. It is also one of the least effective controls in the identity security toolkit.

The structural problems are well understood. Managers are presented with large volumes of entitlement data that they lack the technical context to evaluate meaningfully. Reviews are time-constrained and compete with core business responsibilities. And the path of least resistance, rubber-stamping the existing access profile, carries no immediate consequence, even when it perpetuates risk.

The result is audit-driven apathy: a learned organizational behavior in which reviewers treat certification campaigns as administrative obligations rather than genuine risk management activities. This dynamic is self-reinforcing. Poor review quality produces poor risk detection, which produces false confidence, which produces under-resourced governance programs, which produces more poor review quality.

Meanwhile, the review types most likely to surface meaningful risk, including role entitlement reviews, policy exception reviews, unstructured data access reviews, and non-human identity audits are among the least commonly performed. These reviews require domain expertise and ownership accountability that supervisor-centric models are not designed to capture.

2. The Negligence Gap and its Legal Consequences

2.1 When governance becomes a liability exhibit

For much of the past two decades, the primary consequence of weak identity governance was operational: breaches, data loss, reputational damage. The legal and regulatory exposure, while real, was often addressed through remediation commitments and settlement agreements.

That calculus is changing. SEC disclosure requirements now mandate timely and accurate reporting of material cybersecurity incidents, and regulators are actively evaluating whether organizations exercised reasonable diligence in their security practices, not just whether they complied with a checklist. DORA, effective across EU financial entities, establishes explicit requirements for information and communications technology (ICT) risk management and incident response that extend to third-party access and identity controls. CJIS security policy requirements for criminal justice information access are rigorous and non-negotiable, with access control failures carrying serious federal consequences. For financial institutions and government agencies, auditors are increasingly focused on whether controls operate continuously, not just whether they were in place during the last audit.

In this environment, a post-breach suit or regulatory proceeding will disregard whether an organization had a governance program. It will instead center on whether that program was operational and enforced at the time of the incident. Discovery will seek access logs, certification records, exception approvals, and evidence of policy enforcement. The question will not be philosophical. It will be evidentiary.

An organization that can produce only annual certification records and no evidence of continuous monitoring has not demonstrated compliance. It has demonstrated the gap between its documented policy and its actual security posture, and it has done so in the most damaging possible context.

2.2 Defining the negligence gap

The Negligence Gap is the widening chasm between what an organization's static policy documentation says and what its dynamic identity systems are actually doing between review cycles. It is not a theoretical risk. It is documented exposure that grows larger with every day that passes between governance actions.

The gap has several contributing dimensions:

- First, access drift: the accumulation of entitlements, role assignments (including outdated role definitions, and permission grants that occur between review periods without triggering any governance action.
- Second, detection latency: the elapsed time between an access anomaly occurring and the organization becoming aware of it through its governance process.
- Third, remediation lag: the time between detection and actual revocation or remediation of out-of-policy access.

In a breach scenario, the Negligence Gap is the territory where liability lives. Plaintiff's counsel and regulators will measure it in time: how long did the compromised access exist, when should it have been detected, and what would a reasonable governance program have done differently.

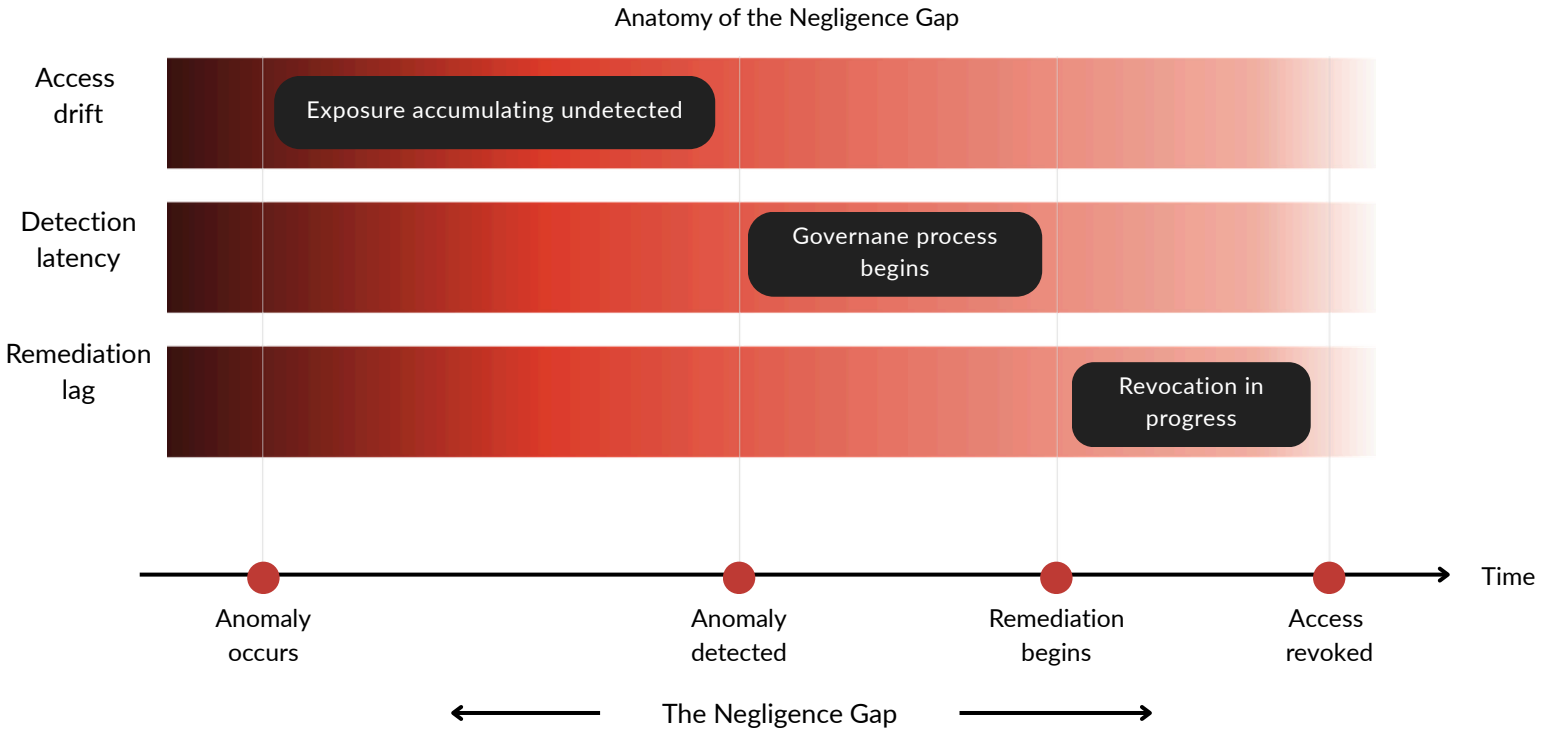

2.3 Non-human identities: the unreviewed attack surface

The identity perimeter has expanded dramatically beyond human users. Service accounts, API tokens, machine identities, RPA bots, and AI agents now constitute a substantial share of total identities in most enterprise environments. These non-human identities often operate with elevated privileges, rarely appear in supervisor review campaigns, and have access scopes that evolve organically as systems and integrations change.

From an adversary's perspective, non-human identities represent a highly attractive target: high privilege, low visibility, and long dwell time before detection. From a legal perspective, compromised service accounts used as lateral movement vectors raise pointed questions about whether organizations maintained effective oversight of all identity types, not just the ones that appear on an org chart.

Governance programs that cannot demonstrate active monitoring and policy enforcement for non-human identities are operating with a structural blind spot that regulators, litigators, and adversaries are increasingly aware of.

3. Active Defense Governance: A Framework for Continuous Control

3.1 The philosophical shift

Active Defense Governance (ADG) represents a fundamental reorientation of how identity governance is conceived and operationalized. The traditional model asks: Who has access? The ADG model asks: Why does this identity have this access, right now, and is that access consistent with current policy and risk posture?

This is not merely a philosophical distinction. It reflects a structural change in how governance functions within the security operations environment. In the traditional model, governance is a periodic function, a campaign that runs, closes, and then waits to resume. In the ADG model, governance is a continuous telemetry stream: every access event, provisioning action, and entitlement change is evaluated against policy in real time, with automated response capabilities that do not require human review cycles to take effect.

The formal access certification campaign does not disappear in this model. It is reframed. Rather than functioning as the primary control mechanism, periodic certifications serve as confirmation and attestation of what continuous monitoring has already detected, managed, and documented. The campaign validates the system and the policies. It is no longer the system.

3.2 Core capabilities of an ADG program

An effective Active Defense Governance program is built on several interdependent capabilities:

Continuous access monitoring

Real-time policy evaluation replacing static review cycles with event-driven governance triggers.

Risk-based certification prioritization

Human review effort directed toward highest-risk entitlements and identities, not distributed uniformly across all access.

Just-in-Time (JIT) access provisioning

Access granted for a defined window and permission scope, with automatic revocation at expiry, eliminating standing privileges as the default state.

Non-human identity (NHI) governance

The same continuous monitoring and policy enforcement applied to service accounts, API credentials, AI, and machine identities that is applied to human users.

Time-to-Revocation (TTR) measurement

Elapsed time from identity change event to access revocation tracked as a primary security metric.

Automated, audit-grade logging

Continuous, tamper-evident records of access decisions, policy evaluations, and governance actions — immediately available for regulatory or legal review.

3.3 Just-in-time governance and Zero Trust alignment

Just-in-Time access governance addresses one of the most persistent structural weaknesses in traditional IGA: standing privilege. In most enterprise environments, access is provisioned and then persists until it is explicitly revoked, which in practice often means that access persists indefinitely. Standing privileges are a primary enabler of lateral movement and privilege escalation in identity-based attacks. JIT governance inverts this model for high risk access requests, not for everyday resources such as email. Rather than treating all privileges as permanent until revoked, it applies time-bound, policy-evaluated grants where standing access poses the greatest risk. Every elevated or sensitive access event becomes a governance decision, not a legacy artifact. The goal is not to create friction for every access request, but to ensure that where risk warrants it, access is deliberate, time-limited, and tied to a verified need rather than standing access that persists simply because it was granted in the past.

This approach aligns directly with Zero Trust architecture principles, specifically the requirement for continuous verification and least-privilege access enforcement. It also produces the governance artifacts that are the most defensible in a regulatory or litigation context: these artifacts are time-stamped, policy-linked, automated, and demonstrate what an organization's governance function does in practice rather than aspires to in theory. When investigators ask why a particular identity had access to a particular resource at a particular moment, a JIT governance model produces a precise, documented answer.

4. Measuring What Actually Matters

4.1 Metrics that reflect security reality

Governance programs have historically been measured by campaign completion rates: the percentage of reviews certified within a defined window. This metric is almost entirely disconnected from security outcomes. A one hundred percent campaign completion rate achieved through rubber-stamp approvals is evidence of administrative compliance, not risk management.

An ADG program requires a different metric set, one oriented toward actual security posture and legal defensibility. Key metrics to consider include:

- Policy enforcement rate (continuous), replacing campaign completion rate as the primary control indicator
- Real-time anomaly detection rate, measuring the percentage of access anomalies detected within defined SLA windows
- Time-to-Revocation (TTR), tracking elapsed time from triggering identity event to access revocation
- Non-human identity coverage, measuring the percentage of non-human identities under active governance
- JIT access adoption rate, tracking the percentage of privileged access granted through JIT mechanisms versus standing privilege
- Audit evidence completeness, measuring the availability and integrity of continuous governance logs for any given period

An ADG program requires a different metric set, one oriented toward actual security posture and legal defensibility.

4.2 Time-to-Revocation as a board-level metric

Time-to-Revocation (TTR) deserves particular emphasis because it directly measures the window of opportunity available to an attacker following a triggering identity event, whether a termination, a role change, an anomaly detection, or a policy violation.

Organizations with mature continuous governance capabilities measure TTR in minutes. Organizations relying on periodic review cycles measure it in days, weeks, or, for access types not covered in standard campaigns, potentially indefinitely (if they measure it at all). This differential is not a nuance. It is the difference between an identity attack that is contained and one that becomes a material breach. For financial institutions and government agencies, that differential matters beyond the security operations center. Examiners increasingly expect a precise, documented answer to how quickly access was revoked following a triggering event, not a reference to the next certification cycle.

TTR is also one of the most effective metrics for board-level reporting because it translates technical governance capabilities into tangible risk reduction outcomes. An organization that can demonstrate an average TTR of under fifteen minutes following a termination event has communicated something meaningful about its security posture, something that a campaign completion rate cannot convey.

Continuous governance program

<15 Min

Average TTR following a termination event or anomaly detection

Access granted for a defined window and permission scope, with automatic revocation at expiry, eliminating standing privileges as the default state.

Periodic review program

Days—or never

TTR for access types not covered in standard campaigns may be indefinite

Revocation depends on the next review cycle. For non-human identities and edge case entitlements, that cycle may never come.

5. Implementation Considerations

5.1 The maturity continuum

The transition from compliance-oriented IGA to ADG is not a single-step change. It is a maturity progression that most organizations will approach incrementally, beginning with the highest-risk identity categories and extending coverage over time.

A practical maturity model for ADG adoption progresses through four stages:

1

Foundational

Centralized identity inventory, consistent provisioning and deprovisioning workflows, and baseline access certification campaigns. This stage is the prerequisite for everything that follows.

2

Risk-Aware

Introduction of risk scoring for identities and entitlements, prioritization of review effort based on risk weight, and initial implementation of anomaly detection capabilities.

3

Continuous

Real-time policy evaluation, event-driven governance triggers, TTR measurement, and integration with security operations tooling. This is the ADG threshold.

4

Defensible

JIT access governance applied where standing access poses the greatest risk, full non-human identity coverage, continuous audit-grade logging, and regulatory reporting automation. This stage represents the litigation-ready posture.

5.2 AI as a force multiplier for continuous governance

The volume of access events, entitlement changes, and identity interactions in a modern enterprise makes manual continuous governance operationally infeasible. AI-driven analytics are not an optional enhancement to an ADG program. They are an architectural requirement.

Applied effectively, AI capabilities within an IGA platform can perform several functions that are beyond the practical reach of human-driven review processes. Behavioral analytics can establish baseline access patterns for each identity and flag deviations that warrant policy evaluation. Machine learning models can assess the risk weight of individual entitlements and access combinations, identifying toxic privilege combinations that standard role reviews miss. Natural language processing can interpret policy documents and map entitlements to policy intent, enabling automated compliance assessment rather than manual interpretation.

Critically, AI-driven access decisions must be backed by verifiable risk scores and documented policy linkages, not simply by model outputs. The governance artifact that matters in a legal or regulatory context is not that an AI made a decision. It is that the decision was made according to a documented, auditable policy evaluation process, with the risk basis recorded at the time of the decision.

Applied effectively, AI capabilities within an IGA platform can perform several functions that are beyond the practical reach of human-driven review processes.

5.3 Repositioning governance as a security function

Perhaps the most significant implementation challenge in transitioning to ADG is organizational rather than technical. Traditional IGA programs are often housed within compliance, HR technology, or IT operations functions. ADG requires that identity governance be treated as a core security capability, integrated with security operations, incident response, and threat intelligence.

Repositioning ADG in this context has practical implications. Governance metrics need to appear in security operations reporting, not only in compliance dashboards. Governance anomalies need to feed into SIEM and SOAR workflows. And governance leadership needs to have a direct line to security leadership, not a reporting relationship mediated through compliance or audit functions.

The shift also requires a change in how governance investments are evaluated. Spending on continuous governance infrastructure is not a compliance cost. It is a security investment with a measurable risk reduction outcome and a direct bearing on cyber insurance premiums, regulatory standing, and litigation posture.

6. The Defensibility Imperative

The central argument of this paper is not that compliance is irrelevant. Frameworks like SOC 2, DORA, and CJIS security policy serve important functions: they establish baseline expectations, provide audit structure, and create accountability mechanisms. The argument is that compliance, as currently practiced by most organizations, is insufficient, and that the insufficiency is no longer merely a security risk. It is a legal and reputational one.

The identity threat landscape has evolved to a point where any governance program that cannot demonstrate continuous, policy-enforced, real-time access control is operating with a structural exposure. Periodic reviews can still play a role, as attestation and confirmation of continuous controls, as a mechanism for policy refinement, and as a regulatory deliverable. But they cannot be the primary control. Risks do not wait for review cycles.

> *The question is not whether your organization has a governance policy. The question is whether you can prove it was working when it mattered and that it continues to work today.*

Organizations that make the transition to ADG are building something that goes beyond a more effective security program. They are building a governance posture that is defensible in the contexts that matter most: a regulatory inquiry, a board conversation, a post-breach investigation, and if it comes to it, a courtroom.

The question is not whether your organization has a governance policy. The question is whether you can prove it was working when it mattered and that it continues to work today.

About this paper

This position paper was developed by the Identity Governance and Administration experts at RSA. It is intended to provide practitioners, program leaders, and executive stakeholders with a framework for evaluating and evolving their identity governance posture beyond periodic compliance activities toward a continuous, defensible Active Defense Governance model.

About RSA

RSA is the identity standard for government agencies, finance, and high-assurance organizations. RSA provides the identity intelligence, authentication, access, governance, and lifecycle capabilities needed to prevent threats, secure access, maintain operations, and surpass compliance. More than 9,000 security-first organizations trust RSA to manage more than 60 million identities across on-premises, hybrid, and multi-cloud environments. For additional information, visit our website to contact sales, find a partner, or learn more about RSA.