SACR

NOVEMBER 2025

MODERNIZING THE IDENTITY STACK: FROM VISIBILITY TO GOVERNANCE THROUGH ENTITLEMENT INTELLIGENCE

FRANCIS ODUM

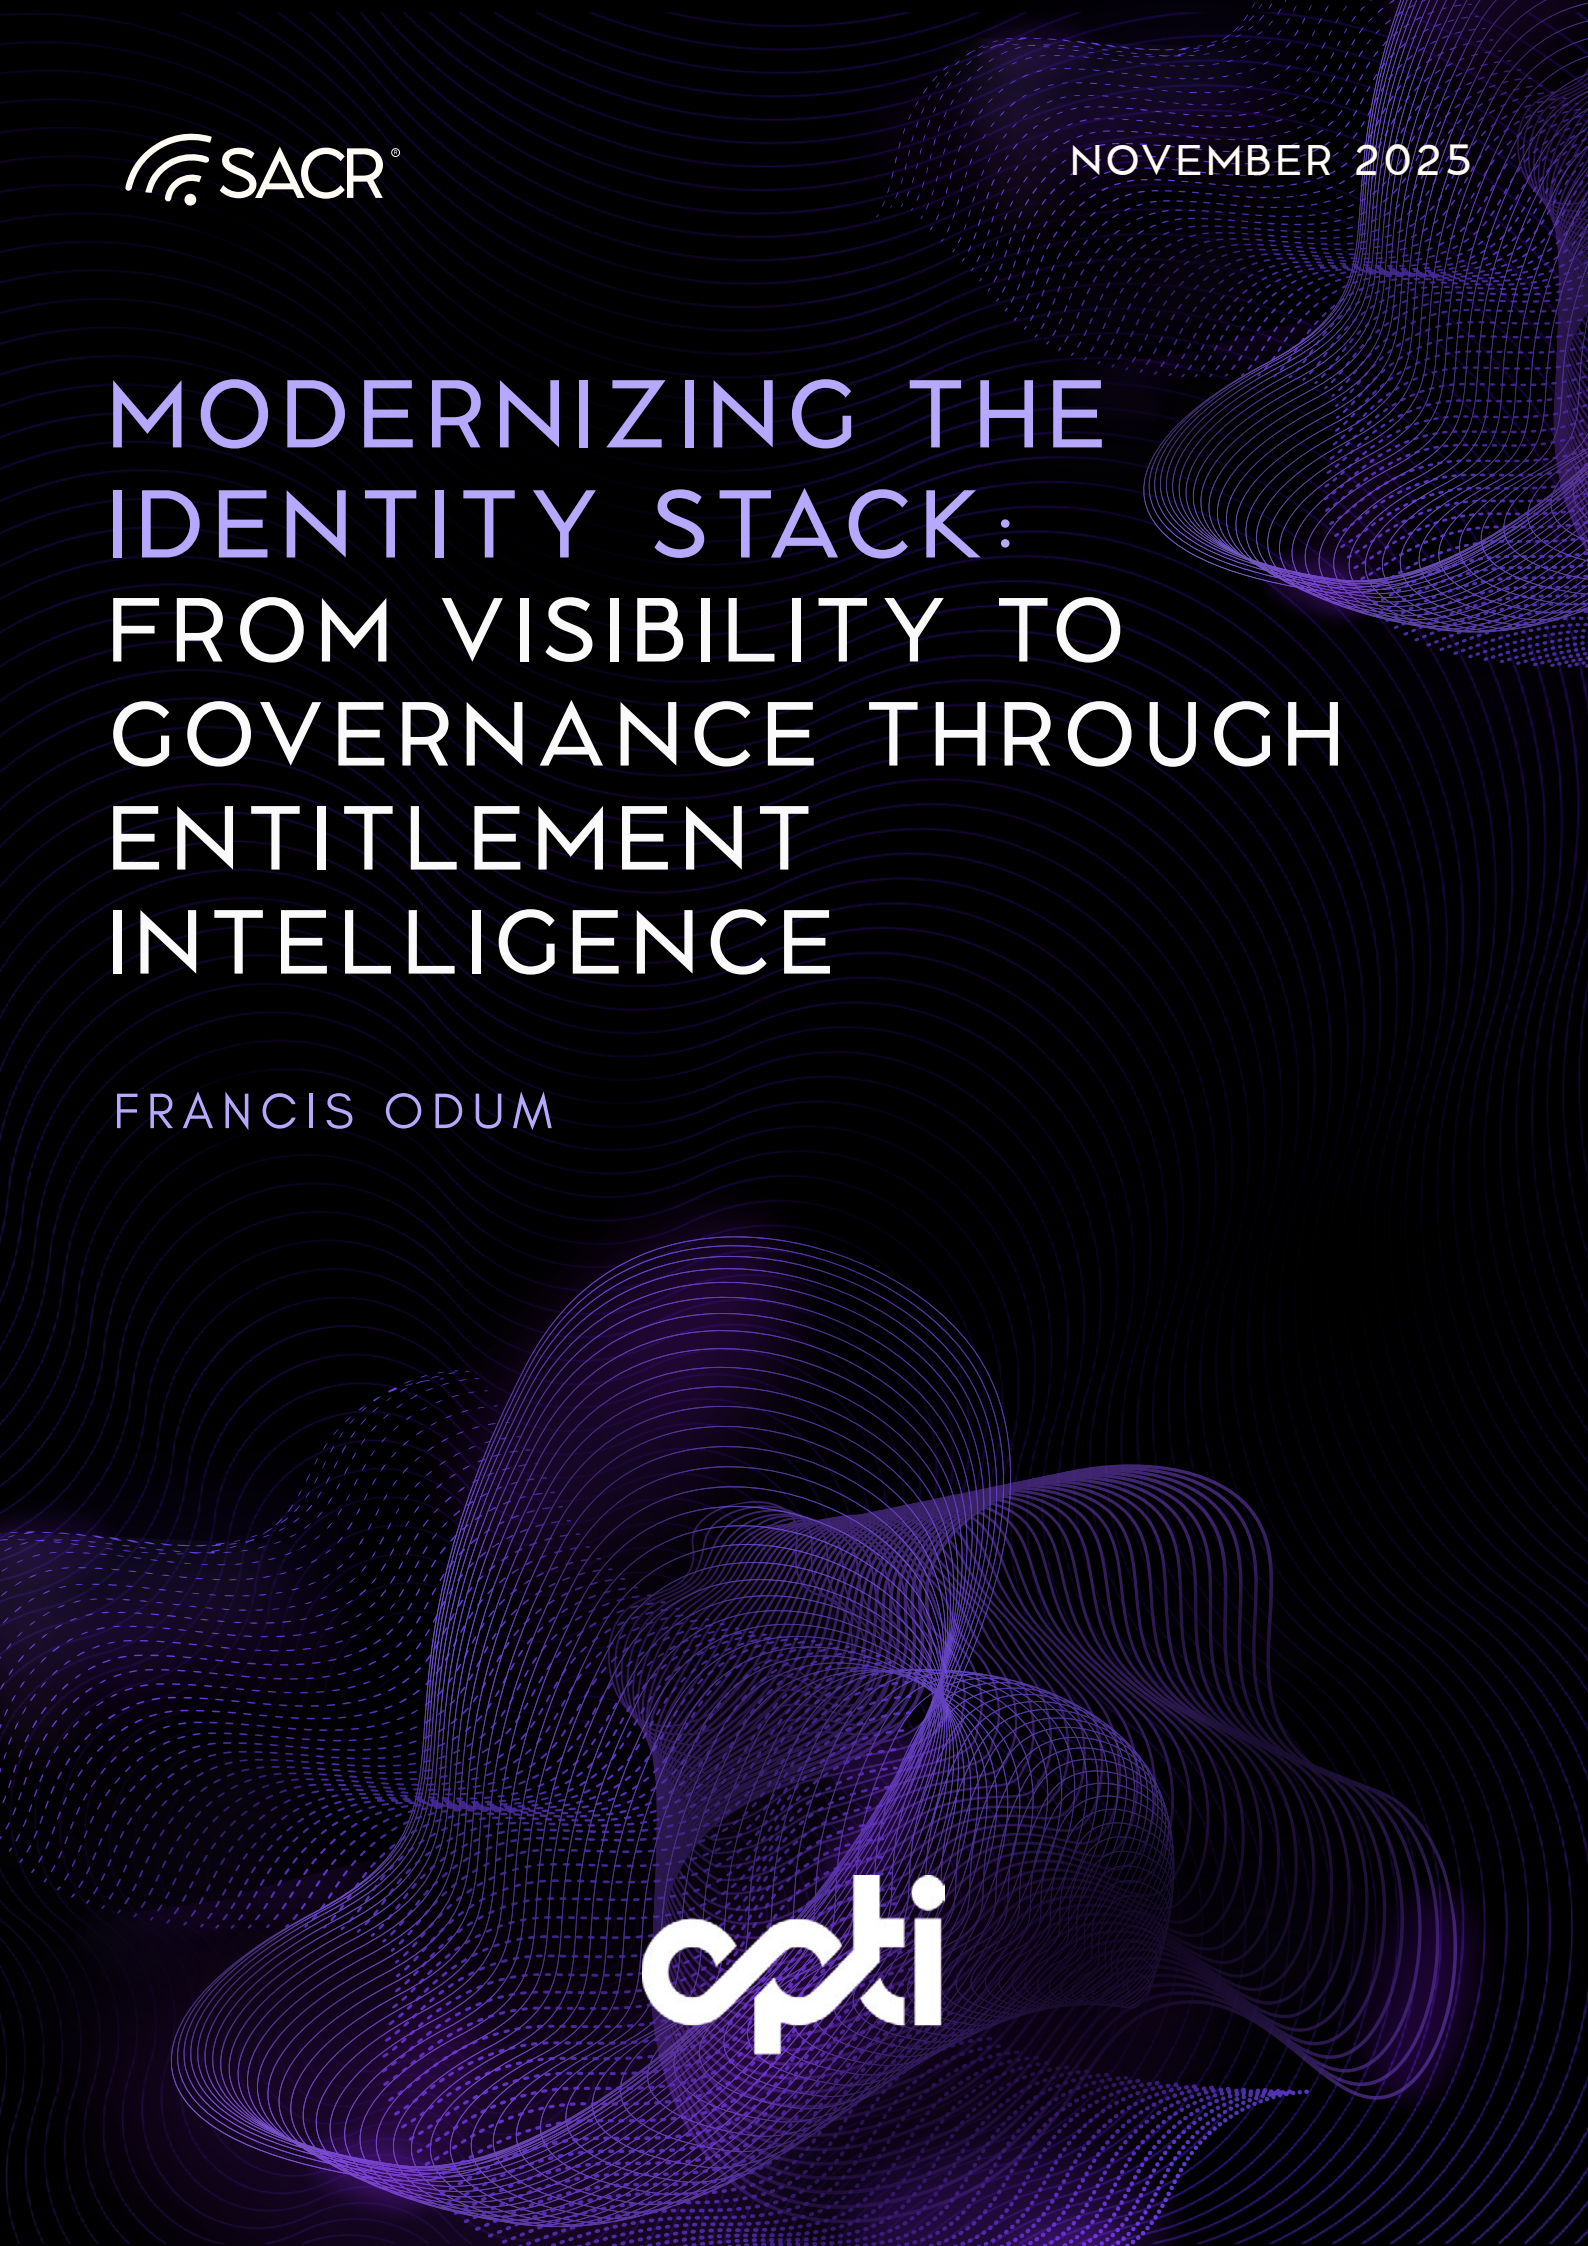

Table of Contents

Research

We explore the newest frontiers of cybersecurity.

Whether you're looking at emerging vendors, evolving threats, or shifting architectures, our timely, opinionated insights help modern security leaders make smarter, faster decisions.

About Software Analyst Cybersecurity Research

SACR is a modern research and advisory firm built for today's cybersecurity leaders. We deliver in-depth, timely analysis across SOC operations, Identity, Network, Cloud, Application Security, Data, and AI Security; equipping CISOs, security teams, founders, investors, and practitioners with the insight they need to navigate high-stakes decisions.

With an engaged community of over **80,000** readers and followers, SACR connects with a global network of cybersecurity decision-makers and innovators. Our access to leaders across categories and industries gives us a direct line to the conversations shaping the market. By pairing these insights with rigorous technical analysis and continuous market tracking, we produce research that is both data-driven and grounded in the realities of modern security operations.

Whether you're seeking clarity on emerging technologies, evaluating vendors, or tracking market shifts, SACR delivers trusted, independent research designed to help you see clearly and decide with confidence.

Author

- **Francis Odum** is the Founder/CEO of the Software Analyst Cyber Research, where he leads the firm's research and engagement with cybersecurity leaders.

Why Entitlements Now Define Identity Risk

- **Credential/identity abuse** continues to dominate modern breach mechanics, and the data is increasingly consistent across sources. Verizon's 2025 DBIR shows that the use of stolen credentials remains the most common initial action in breaches, accounting for 31% of incidents, while the human element is involved in 76% of breaches overall. Despite years of investment in authentication controls, attackers continue to succeed by inheriting or abusing legitimate access.
- **Breach costs -** Global data breach costs in 2025 averaged USD 4.44 million, while U.S. breaches surged to USD 10.22 million, with 97% of AI-related incidents lacking proper access controls (IBM Cost of data breach 2025)
- **Stolen/compromised credentials** were the most common initial attack vector in IBM's 2025 Cost of a Data Breach (16% of breaches) and *take the longest to identify and contain* (~8-10 months) (IBM Newsroom) and (2025 IBM report)
- **The scale of identity attacks is staggering**. Microsoft reports ~600M identity attacks/day; Entra data shows >7,000 password attacks/second and that *password-based attacks account for 99%* of daily identity attacks. (Microsoft)
- **Initial access mix is shifting but creds stay critical.** Mandiant's M-Trends shows *stolen credentials* rising to the #2 initial vector (16%), behind exploits.

In this report, we focus on the challenges organizations face when managing Identity governance and administration (IGA) and entitlements. As identity environments expand across cloud, SaaS, and infrastructure, the gap between visibility and governance continues to widen. Entitlements sit at the center of this problem.

Identity Entitlement Model: The Missing Authorization Layer

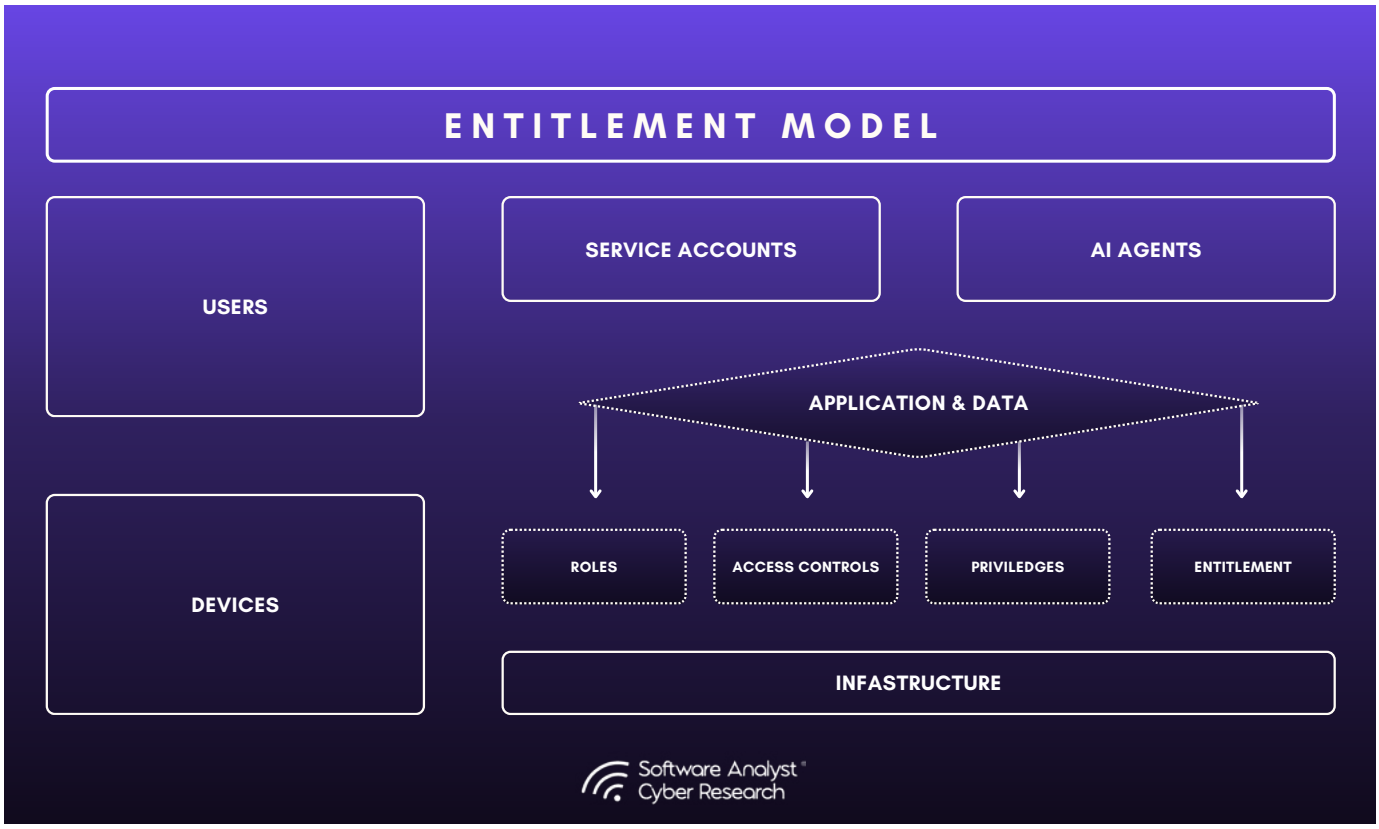

In the identity security market, an entitlement model is a structured framework that defines who has access to what, and more importantly, why. It acts as the logic layer that maps identities such as users, service accounts, and devices to resources including applications, data, and infrastructure through roles, permissions, and privileges.

At its core, an entitlement represents a discrete access right. This could be read access to an S3 bucket, an administrator role in Salesforce, or write privileges to a GitHub repository. The entitlement model is how an organization defines, groups, and manages these permissions in a consistent way across systems.

Tying It Together

In a typical enterprise environment, identity responsibilities are fragmented. Okta authenticates the user. SailPoint tracks what entitlements that user should have based on HR attributes and lifecycle data. AWS IAM ultimately enforces what entitlements the user actually has within cloud infrastructure.

A strong entitlement model sits beneath these systems. It maps and reconciles the differences between intended access and actual access, flags risk conditions, and supports policy enforcement. Rather than replacing IAM or IGA, the entitlement model provides the connective tissue between them.

When implemented effectively, the entitlement model defines and governs who can do what across disparate systems. Strong entitlement management platforms provide deep visibility into permissions, automate entitlement lifecycles, and enforce policy consistently. Over time, this model becomes the authorization backbone of the identity stack, underpinning IGA, CIEM, PAM, and Zero Trust initiatives.

Modern IGA Challenges Faced By CISOs

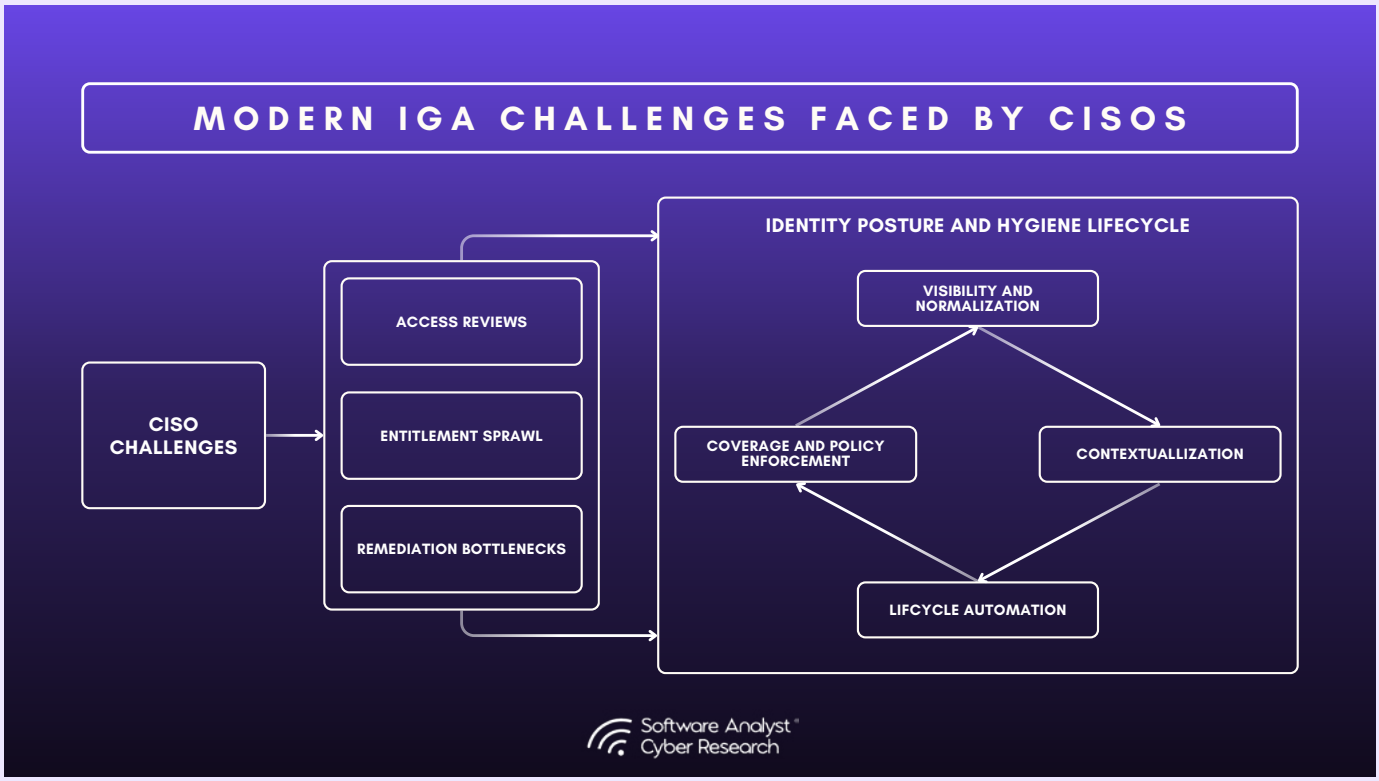

IGA Access Reviews Have Become Administrative Rituals Rather Than Meaningful Controls

Access certification campaigns, once the backbone of enterprise identity governance, have steadily deteriorated into highly manual, compliance-driven exercises that generate limited security value. Most organizations now run quarterly or annual reviews covering hundreds of applications and tens of thousands of entitlements. Yet reviewers are typically presented with raw, technical permissions that lack business context, usage data, risk indicators, or peer comparison. Unsurprisingly, these reviews result in near-universal approval, regardless of whether the access is appropriate.

Security leaders increasingly acknowledge that access reviews, in their current form, have become rubber-stamping exercises. Review fatigue is widespread. Managers are often asked to certify hundreds of items with little understanding of what those permissions actually enable. Gartner has repeatedly highlighted this dynamic, noting that excessive approval volumes, insufficient context, and over-complex role models lead to perfunctory reviews that fail to reduce risk. The net effect is that entitlements, once granted, often persist indefinitely.

Empirical data reinforces this gap. In a 2024 survey conducted by SailPoint, 58% of identity leaders reported that their access reviews were ineffective, while 42% stated that reviewers lacked the context required to make informed decisions. When paired with IBM's findings that credential misuse remains a leading breach vector, the limitation becomes clear. Organizations cannot rely on traditional IGA campaigns as their primary control for identity-related risk. In a report on Scatted Spider, SACR has also called out the importance of identity intelligence and visibility as a key aspect of defense of these emerging attacks.

Entitlement Sprawl has Outpaced the Governance Capabilities of Most Enterprises

Across large organizations, entitlement data has grown faster than any other identity domain. Enterprises now operate in environments where thousands of SaaS applications, expanding multi-cloud footprints, and a growing population of machine identities collectively generate millions of discrete entitlements. Each system introduces its own permission taxonomy, authorization logic, and privilege hierarchy.

The result is what many security leaders describe as entitlement sprawl: a proliferation of fine-grained, inconsistently defined permissions that traditional IGA platforms were not designed to model or govern at scale.

This fragmentation undermines the core value proposition of governance. When IGA platforms lack a normalized and contextual view of entitlements, access certifications become superficial, role engineering becomes unmanageable, and policy enforcement becomes inconsistent or incomplete. Organizations attempt to compensate through manual spreadsheet mapping, custom connectors, or services-heavy engagements, but these approaches increase cost without establishing durable visibility.

The operational risks are well documented. Verizon's 2025 Data Breach Investigations Report attributes roughly **60%** of breaches to the human element, with credential misuse and third-party exposure remaining dominant attack paths. IBM's 2025 Cost of a Data Breach report further notes that incidents involving compromised credentials take more than eight months to fully contain, largely because over-privileged access expands the blast radius of compromise. While AI-assisted tooling is beginning to improve response times, the underlying problem is consistent across these findings. Authentication is no longer the bottleneck. Authorization sprawl is.

Remediation Bottlenecks Prevent Identity Programs from Reducing Actual Risk

Even when organizations successfully identify over-entitled users, toxic combinations, or noncompliant access paths, remediation is often slow, fragmented, and operationally risky. This is perhaps the most overlooked challenge in IGA programs: although the gap between *knowing* and *fixing is improving, it's still far too lengthy to deal with real-time threats posed by AI agents*.

Most IGA tools were designed to orchestrate approvals, not to implement precise, dependency-aware entitlement changes across SaaS, cloud, and legacy environments. As a result, remediation commonly devolves into IT tickets that require coordination across identity teams, app owners, and security staff. Each change must be risk-assessed, tested, approved, and sequenced correctly especially in applications where permissions are deeply intertwined with business workflows.

This operational friction leads to persistent remediation backlogs. A 2025 survey from Identity Defined Security Alliance found that **33%** of organizations believed they were unable to remediate identified identity risks quickly enough, citing coordination and tooling limitations as primary obstacles. Even entitlements flagged as unused or excessive are commonly deferred due to fear of business disruption, unclear ownership, or lack of automated rollback mechanisms.

The consequence is that excessive privileges, orphaned accounts, and misaligned roles remain in place long after they are identified. This creates a sustained risk window in which attackers can exploit dormant or poorly understood access. Mandiant's M-Trends 2024 report underscores this reality, showing that incidents involving unauthorized credential use often exhibit longer attacker dwell times due to delayed or incomplete privilege cleanup.

Market Response to Entitlement Complexity

Identity Security Posture Management (ISPM)

Identity Security Posture Management (ISPM) focuses on continuously assessing and improving the security configuration of identity systems across cloud, SaaS, and IAM platforms. Its value lies in revealing misconfigurations, unused privileges, weak MFA enforcement, stale accounts, insecure OAuth grants, and toxic permission combinations that accumulate as identity ecosystems expand. ISPM tools serve as the "exposure radar" for the identity stack, surfacing risk conditions that IGA platforms traditionally miss because IGA operates on static, periodic certification cycles rather than real-time posture telemetry. While ISPM does not replace IGA, it amplifies its effectiveness by feeding governance programs with accurate, high-tempo visibility of misaligned entitlements, privilege drift, and policy violations. In essence, ISPM defines *where* risk exists; IGA defines *what must be done* about it; the entitlement model provides the semantic layer needed to reconcile these signals across systems.

Light-IGA and Identity Access Reviews

Light-IGA solutions are typically modern, SaaS-native platforms or access-review focused tools designed to improve governance workflows without the role engineering overhead and implementation complexity of legacy IGA suites. These platforms emphasize faster onboarding, intuitive reviewer experiences, streamlined access reviews, and partial automation of entitlement lifecycles. Their growth reflects a reality many organizations face. Traditional IGA programs are difficult to sustain and frequently degrade into low-value, rubber-stamped reviews. Light-IGA improves the mechanics of governance, but its long-term effectiveness still depends on the quality and context of the entitlements being reviewed. This is where the entitlement model becomes critical. Without normalized and meaningful entitlement intelligence, such as understanding what a permission actually enables, whether it is actively used, or how it compares to peer access, even modern Light-IGA solutions risk reproducing the same governance fatigue as earlier tools.

In practice, Light-IGA delivers the workflow. The entitlement model delivers the insight. IGA provides the compliance structure that binds them together.

What a Modern Entitlement Management Platform Must Deliver

To address these challenges, a strong entitlement management platform provides continuous, automated, and context-aware control over who has access to what, and why. The proliferation of SaaS platforms, cloud infrastructure, and non-human identities such as service accounts, pipelines, and AI agents has pushed entitlement volume beyond the limits of manual governance. Without automation and intelligence, entitlement sprawl, privilege creep, and compliance exposure are unavoidable.

Effective platforms embed unified visibility, automation, and policy enforcement throughout the identity lifecycle. This allows organizations to move away from reactive, campaign-based governance toward continuous control that adapts to change in near real time.

Visibility and Normalization

A robust entitlement management platform provides comprehensive visibility and normalization across cloud, on-premises, and disconnected systems. This unified view is foundational for identifying excessive, unused, or risky access, which remains a common root cause of breaches and audit failures. By ingesting entitlements from hundreds of systems into a single queryable model, organizations can analyze AWS IAM roles, Azure AD groups, and Okta assignments under a common schema.

Contextualization

Ideally, these solutions can attach business context: mapping entitlements to departments, risk levels, compliance domains (SOX, HIPAA), or job functions. This transforms raw permissions into meaningful insights, a key differentiator in governance and compliance.

Lifecycle Automation

They can automate the entire identity lifecycle from onboarding, role changes, to offboarding, so that access is provisioned and deprovisioned quickly and accurately, eliminating orphaned accounts and reducing manual workload. Automation is critical for handling the sheer volume and velocity of access changes in today's dynamic environments, and for providing timely evidence of least-privilege enforcement to auditors. They automate entitlement provisioning and deprovisioning, triggered by HR or identity lifecycle events. Strong systems enforce least privilege and separation of duties (SoD) automatically.

Risk Reviews, Access Reviews and Governance

A modern entitlement management platform enforces granular, risk-based policies such as least privilege and separation of duties, using AI or advanced policy engines to detect toxic combinations and anomalous access. This approach not only strengthens security but also meets regulatory requirements for provable, continuous access remediation. They provide powerful access review, attestation, and anomaly detection capabilities. This helps organizations satisfy identity governance and administration (IGA) requirements.

They transform access reviews from manual exercises into meaningful, context-rich validations. By surfacing what an entitlement actually grants, how it is used, and how it compares to peers, reviewers can make informed decisions, reducing unnecessary access and audit fatigue

Coverage and Policy Enforcement

A good solution extends governance to all identities: human and non-human across the entire application landscape, consolidating identity data and automating controls even for legacy or custom systems that lack modern APIs. This comprehensive reach eliminates blind spots and ensures that all access is governed, not just the easy-to-integrate systems.

Modern platforms embed AI or policy engines to detect toxic combinations, unused entitlements, or violations of least privilege. They can recommend or auto-remediate changes, where appropriate.

Strategic Guidance for CISOs

Prioritize Entitlement Intelligence as a Foundational Layer

CISO's should explicitly budget for and implement a dedicated Entitlement Intelligence platform (like Opti) that acts as an "authorization brain." The key rationale for this is **authorization sprawl** and not authentication, has become the primary identity risk. Traditional Identity Governance and Administration (IGA) and Identity and Access Management (IAM) tools are insufficient because they lack a semantic understanding to normalize and contextualize permissions across modern, hybrid environments.

Modernize IGA Reviews to a Risk-Based, Continuous Process

Organizations should end, or significantly restructure, traditional quarterly and annual access certification campaigns that have degenerated into administrative rituals. Instead, identity governance must shift toward risk-driven, continuous review models.

Modern approaches emphasize reviewer context and decision quality. Agentic reviewer experiences, peer baselines, entitlement usage patterns, risk classifications, and policy explanations transform access reviews from mechanical approvals into meaningful control points. The objective is not to review everything on a fixed schedule, but to focus governance effort where risk is highest and change is occurring. This transition is essential if access reviews are to reduce real risk rather than simply satisfy compliance checkboxes.

Proactively Address Non-Human Identity Governance

CISO's must also proactively develop a strategy to map, monitor, and enforce guardrails for entitlements granted to non-human identities, including service accounts and emerging AI agents. Preferring solutions that deliver the ability to extend entitlement mapping to non-human identities directly addresses a growing security concern. Delaying this will leave a significant attack vector exposed as automation increases. This is against the backdrop of a large emergence of AI in the identity ecosystem (see "From Complexity to Control: Using AI and Automation to Transform Enterprise Identity Security").

Opti

Opti is an emerging identity security vendor seeking to leverage LLM advancements in order to redefine how enterprises understand, govern, and remediate entitlements across complex hybrid environments. Its platform is positioned not as a traditional IGA tool, but as an AI-native entitlement intelligence and governance layer that acts as the "authorization brain" for the broader identity ecosystem. Opti's core value proposition is driven by specialized fine-tuned LLMs and centers on unifying fragmented permissions data, providing context-rich and actionable insights, and automating previously manual governance processes through dependency-aware workflows.

This approach targets a persistent gap in modern identity architectures. While IAM and IGA platforms authenticate users and manage workflows, they generally lack a deep understanding of what granted access enables or how misaligned entitlements should be corrected at scale. Opti's aim is to close this gap by introducing a semantic layer that makes entitlement meaning operational.

Opti's central thesis is that entitlement complexity (not authentication) is now the primary identity risk in modern enterprises. The vendor asserts that existing IGA platforms struggle because they ingest raw, unstructured permissions from hundreds of systems without the ability to normalize, contextualize, or reason about them. Opti instead proposes an entitlement intelligence layer that acts as a shared foundation for governance, operations, and risk reduction.

At the core of this approach is a proprietary **entitlement model**, expressed as a triplet:

- **Identity (User/Group) + Resource + Permission**

Unlike traditional models that treat entitlements as strings or role assignments, Opti attempts to understand the *meaning* and *business impact* of each permission. For example, "Billing Admin" surfaces across Salesforce, Zoom, and NetSuite, yet implies fundamentally different scopes and privileges in each. Opti's architecture is designed to resolve these differences by embedding permissions into a normalized semantic graph.

This normalization serves as the basis for policy evaluation, risk scoring, access recommendations, and remediation planning, effectively enabling Opti to function as a real-time authorization intelligence engine for identity teams.

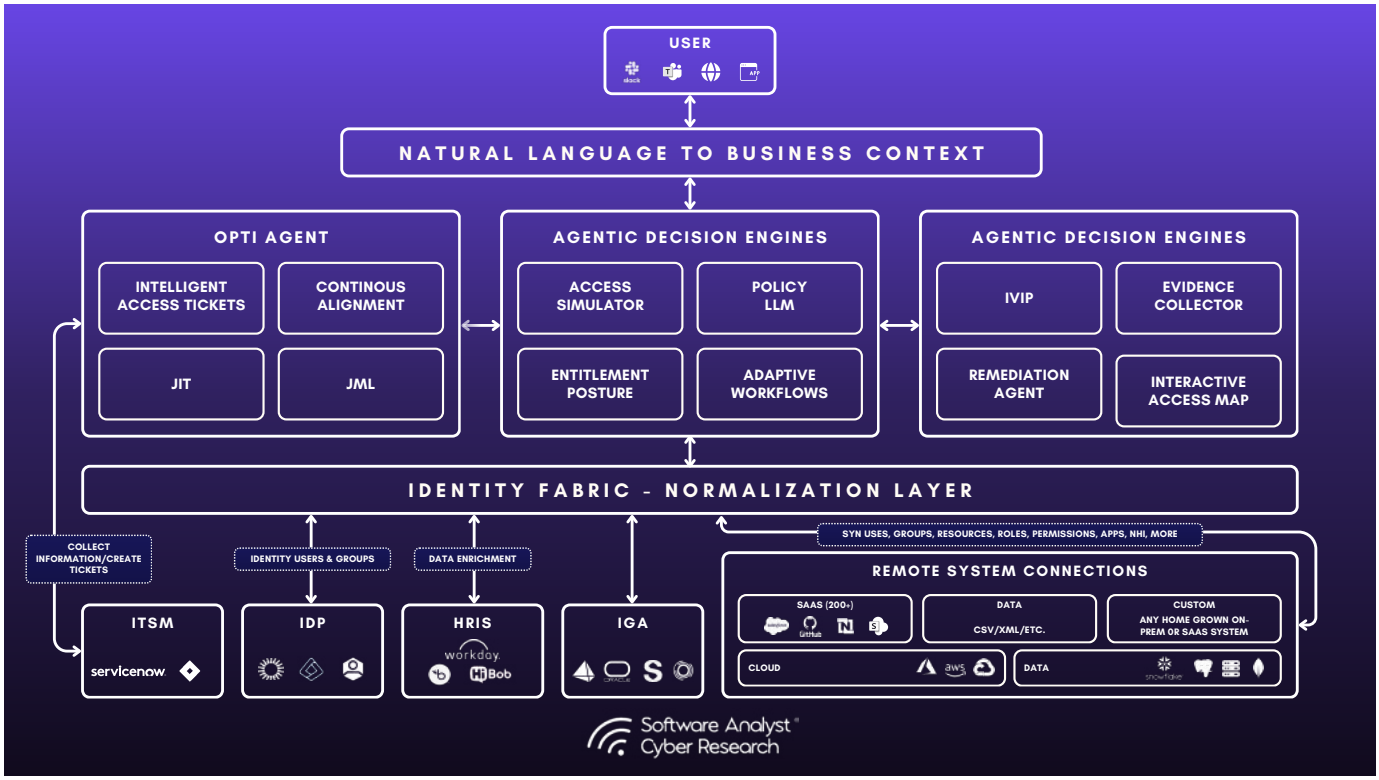

Architectural Approach: Graph-Based Entitlement Ontology

Opti claims to have built a specialized, multi-model architecture leveraging fine-tuned LLMs, predominantly based on Mistral open-source models. These models are trained not on general language tasks but on specific identity and entitlement semantics, enabling capabilities such as:

Technical Architecture and Agentic AI

Opti's architecture is built for scale and flexibility:

1. **AI-Driven Entitlement Model:** Fine-tuned LLMs trained specifically for identity and access management tasks, supported by a dedicated data science team. These models power entitlement normalization, natural language understanding, and policy evaluation
2. **Graph Database Backend**: Enables complex relationship mapping, attack path visualization, and cross-system analysis, underpinning the platform's access intelligence and risk remediation capabilities
3. **Modular AI Agents:** Specialized models that handle different functions: entitlement understanding, workflow modeling, policy translation invoking LLMs only when needed to ensure performance at scale
4. **Natural Language Policy Engine:** Policies and exceptions can be defined and evaluated in natural language, enabling intuitive rule creation and compliance mapping
5. **Rapid Integration:** Opti claims rapid integration even with legacy or homegrown systems, using APIs or headless browser automation, reducing the "professional services tax" that plagues traditional IGA deployments

The platform anchors these models in a GraphDB-based entitlement map, which stores normalized entitlements, identity relationships, and usage baselines. LLMs operate as reasoning layers on top of the graph. Importantly, Opti states that all model outputs are compiled into deterministic, sandboxed code for evaluation, ensuring predictable and auditable control outcomes rather than opaque AI inferences. This "LLM-assisted, rule-enforced" pipeline is designed to mitigate concerns about hallucinations or unpredictable behavior common in generative AI systems.

Opti Platform Features

Detecting Excessive Access and Dormant Accounts

Opti identifies overprivileged users, dormant identities, and entitlement misalignments using AI-driven analytics and policy evaluation. The platform flags unused, risky, or excessive entitlements and can surface toxic combinations such as SoD violations that are difficult to detect across fragmented identity environments.

Streamlining and Contextualizing Access Requests

Access requests are processed through a policy and peer-comparison engine. When a user requests access, Opti analyzes the justification, compares peer entitlements, and recommends the least-privileged, just-in-time, or time-bound alternative. These recommendations are explainable and auditable, supporting compliance and reducing risk. Users can request access in natural language, (for eg; “I need access to the US payroll dashboard for Q4 reporting”), and the system interprets intent, maps it to the correct entitlements, and enforces guardrails.

Automating User Access Reviews and Certifications

Opti transforms periodic user access reviews from a manual checkbox exercise into a continuous, risk-based certification process. Review campaigns are automatically scoped and prioritized based on risk signals, peer comparison, and actual usage, so reviewers see what truly matters instead of blindly re-approving stale access. For each entitlement, Opti provides an AI-generated summary of what it grants, who else has it, and how it’s being used, alongside clear approve/revoke recommendations and impact analysis. All reviewer decisions and rationales are captured as structured, auditable evidence, giving auditors end-to-end traceability while dramatically reducing “rubber-stamping” and review fatigue.

Access management intelligence:

Opti delivers access management intelligence with features such as peer comparison analysis, right-sizing of permissions (scope and time-based/JIT), and AI-driven access request recommendations. The system’s analytics provide reviewers with context: what access an entitlement grants, what peers have, and how usage patterns compare making access reviews meaningful and reducing “rubber-stamping.

Determining and Mapping Entitlement Changes (Including for AI Agents)

While this capability is still in its early inception, Opti’s entitlement mapping is dynamic and extends to both human and non-human identities, including service accounts and, prospectively, AI agents. The platform’s AI models can map, monitor, and enforce guardrails for entitlements granted to AI agents, a growing concern as organizations automate more business processes. The system can auto-revoke excessive or risky entitlements, ensuring that only the minimum required permissions are granted at any time

Automating Access and Ensuring Continued Compliance

Opti automates the full lifecycle of entitlement provisioning and deprovisioning, triggered by HR or identity events, enforcing least privilege and separation of duties (SoD) by default. The platform’s AI models ingest and normalize entitlements from hundreds of systems: cloud, SaaS, on-premises, and legacy into a single, queryable identity fabric. This normalization is not just syntactic but semantic, enabling the system to understand the business and risk context of each entitlement.

Differentiators and Market Impact

1. Proprietary Models for Semantic Normalization at Scale
Where legacy IGA systems depend heavily on human-defined roles and brittle connectors, Opti aims to infer entitlement structure automatically across SaaS and on-premises systems. If this approach holds at enterprise scale, it could materially reduce onboarding effort and shorten time to value in large, heterogeneous identity environments where entitlement models are often fragmented and incomplete.

2. AI-Driven Remediation Planning
Opti's remediation design represents a meaningful departure from posture-focused identity tools that stop at detection. The platform attempts to generate dependency-aware remediation plans that account for identity provider group structures, application-level dependencies, and change sequencing. If consistently reliable, this capability could address a long-standing operational gap in IGA programs

3. Natural Language Policy-to-Action
By compiling natural-language policies into deterministic, guard-railed graph queries, Opti can explain findings in the language of auditors and risk teams. This linkage from policy → evaluation → finding → remediation remains a weakness across IGA and CIEM markets.

4. Agentic Reviewer Experience
Opti emphasizes reducing "rubber-stamp" governance by augmenting reviews with peer baselines, usage context, risk classifications, and policy explanations. Given widespread reviewer fatigue in traditional access certification processes, this enhancement resonates strongly with organizations seeking more effective governance outcomes.

5. Identity Fabric With Learned Integrations
Opti claims its platform can scale integrations not by manually coding each connector but by learning entitlement structures using its multi-model approach. While this claim requires validation, the underlying direction aligns with a broader industry trend to reduce connector dependency and move toward metadata-driven interoperability.

Use Cases and Fit in the Identity Ecosystem

Opti's Strategic Wedge

The IAM and IGA market in 2025 remains large, fragmented, and under increasing pressure from credential-based attacks, regulatory scrutiny, and entitlement sprawl across SaaS and cloud infrastructure. Established platforms such as SailPoint, Okta, CyberArk, and Microsoft remain deeply embedded in enterprise environments but often require long implementation cycles and struggle to adapt governance models to modern authorization complexity.

Opti's go-to-market strategy is best understood as a wedge rather than a direct replacement play. The platform enters through visibility and risk remediation capabilities, then expands into access intelligence, with an eventual ambition to displace portions of traditional IGA workflows. Early adopters view this as a calculated alternative to scaling legacy IGA, provided the platform can sustain enterprise performance and operational safety.

In practice, Opti is complementary to existing identity providers and IGA systems today. This positioning reduces switching friction while allowing organizations to address high-impact risk areas without overhauling their identity stack.

While Opti is not a full IGA suite at present, it occupies a strategically valuable layer between posture visibility tools(ISPM/CIEM) and governance orchestration platforms (IGA/PAM). . Its extension into non-human identity governance may further broaden its relevance as automation and AI agents proliferate.

View Opti as a Strategic Wedge, Not an Immediate IGA Replacement

A practical adoption model is to deploy Opti as an intelligence and remediation layer between existing identity providers (IDP/IAM) and IGA platforms. This approach targets the most acute gaps in visibility and authorization risk(ISPM/IVIP) without forcing wholesale replacement. While Opti's longer-term trajectory points toward broader governance coverage, its near-term value lies in entitlement intelligence and controlled remediation.

Architecturally, the platform reflects a shift toward AI-defined authorization governance, addressing areas where incumbent IAM and IGA platforms continue to struggle.

Risks and Considerations

As with any emerging platform, buyers should temper ambition with validation:

1. **Scalability and connector claims** should be tested on the buyer's most complex systems (e.g., SAP, NetSuite, homegrown apps).
2. **AI explainability and audit alignment** should be reviewed with compliance stakeholders.
3. **Remediation safety** especially for high-impact entitlements requires evidence of rollback, simulation, and change-testing capabilities
4. **Breadth of application mappings** needs ongoing expansion to sustain the normalization advantage.

Analyst Bottom Line

Opti is still relatively new into the category. The critical question they need to face is execution: If Opti can deliver enterprise-scale performance and prove its operational and ROI claims, it has the potential to redefine the identity security landscape for the AI era. Its agentic AI capabilities: semantic entitlement mapping, natural language intelligence, user-first remediation, and rapid integration are not just incremental improvements but foundational advances that directly address the most pressing challenges facing modern identity programs.

If successful, Opti could become the reference architecture for AI-driven identity security, as the market consolidates toward unified, intelligent platforms that bridge the gaps between visibility, governance, and action. Opti represents a new class of identity security vendor centered on entitlement intelligence, a discipline increasingly viewed as foundational to modern access governance and Zero Trust maturity. Its multi-model AI architecture, graph-based normalization, and policy-to-remediation pipeline differentiate it from both legacy IGA tools and posture-only visibility platforms. If Opti can demonstrate consistent accuracy and operational safety at enterprise scale, it has the potential to become a strategic layer in the identity security stack, influencing how organizations define, govern, and maintain least privilege across cloud and SaaS ecosystems.

Software Analyst®
Cyber Research

business

personal

softwareanalyst.io

SACR®

Trusted research. Sharp insights. Real conversation.

CISO	VENDOR
SECURITY TEAMS	INVESTORS