IDEAS: Uniuyo Journal of Philosophy
and Multi-Disciplinary Studies
ISSN: 3093-009x
Vol. 1, No. 6, June 2026
ideasjournaluniuyo@gmail.com
www.ideasuniuyojournal.com

Harmonization of User Privileges on Multi-Cloud Platforms Using Python: The Sentinel Harmonization Engine

Udeagwu Chinedu Kelvin (U22CYS1126)
Department of Cyber Security, Faculty of Computing
Air Force Institute of Technology (AFIT), Kaduna, Nigeria
Email: udeagwuchinedu579@gmail.com

Abstract

The adoption of multi-cloud strategies across enterprise organizations has introduced a critical security challenge: which is the inconsistent management of user privileges across heterogeneous cloud platforms. Each major cloud provider as Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP) implements Identity and Access Management (IAM) under entirely different architectures, resulting in privilege fragmentation, privilege creep, and widening security gaps that existing platform-native tools cannot adequately address. This study presented the Sentinel Harmonization Engine, a Python-based system designed to ingest IAM data from all three platforms, normalize user privileges into a unified HIGH, MEDIUM, and LOW classification scale, detect cross-platform inconsistencies through an automated harmonization algorithm, and generate actionable remediation recommendations. Testing on a simulated multi-cloud dataset identified seven distinct privilege inconsistencies and produced a Security Health Score of 50 out of 100, indicating a critical risk environment. The system is delivered through an interactive Streamlit dashboard with support for CSV and Excel report exports, email alert simulation, and scan history logging. The findings of this study established that automated privilege harmonization across multi-cloud environments is technically achievable using Python, and the Sentinel Harmonization Engine provides a practical, extensible foundation for enterprise-grade multi-cloud Identity and Access Management governance.

Keywords: Multi-Cloud Security, Identity and Access Management (IAM), Privilege Harmonization, Least Privilege Principle, Python Automation.

Introduction

The digital transformation of public and private sector organizations worldwide has catalyzed an unprecedented shift from traditional on-premises infrastructure to cloud-based computing models. Enterprises today no longer rely on a single cloud provider. Instead, they strategically distribute workloads across multiple providers a practice known as a multi-cloud strategy to achieve greater scalability, resilience, cost efficiency, regulatory compliance, and freedom from vendor lock-in (Ali, Khan, & Vasilakos, 2015). While this approach delivers significant operational benefits, it simultaneously introduces one of the most complex and underappreciated security challenges of the modern era: which is the inconsistent governance of user identities and access privileges across cloud environments.

At the heart of cloud security lies Identity and Access Management (IAM) the framework that determines who can access which resources, under what conditions, and with what level of authority. The principle of least privilege, a cornerstone of IAM governance, demands that every user be granted only the minimum level of access necessary to perform their designated function. In a single-cloud environment, enforcing this principle is already non-trivial. Across multiple cloud providers, it becomes substantially more complex because each provider implements IAM through a fundamentally different architecture.

AWS governs access through JSON-based policies attached to users, groups, and roles. Azure integrates Role-Based Access Control (RBAC) with Azure Active Directory, tying permissions to predefined or custom roles assigned at hierarchical scopes. GCP employs a resource hierarchy model where roles are bound to members at the project, folder, or organization level. The structural incompatibility between these three models means that a single organizational user say, a system administrator may hold dramatically different privilege levels across the three platforms, without any unified mechanism to detect, compare, or reconcile those differences.

This research addresses that problem directly. as it presents the design and implementation of the Sentinel Harmonization Engine: a Python-based prototype system capable of ingesting IAM configuration data from AWS, Azure, and GCP; normalizing the data into a unified representational model; automatically detecting privilege inconsistencies; computing risk scores; and generating remediation recommendations all through an interactive web dashboard. The system represents a move from theoretical frameworks to practical, deployable tooling in the field of multi-cloud privilege governance.

Cloud computing fundamentally transformed organizational information systems by offering on-demand access to computing resources including Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a

Service (SaaS). The ability to scale operations without proportional capital investment in physical infrastructure made cloud adoption irresistible to enterprises across every sector from financial services and telecommunications to government, education, and healthcare.

The evolution from single-cloud to multi-cloud deployment is a natural progression of this trend. Organizations operating across AWS, Azure, and GCP simultaneously gain the ability to distribute workloads based on each provider's comparative strengths, meet regional data sovereignty requirements, ensure business continuity through platform redundancy, and negotiate more competitive pricing. According to Amajuoyi, Nwobodo, and Adegbola (2024), multi-cloud strategies have become particularly prevalent among businesses, governments, financial institutions, and educational organizations globally, making them a defining feature of contemporary enterprise IT architecture. However, the same heterogeneity that makes multi-cloud deployment advantageous also makes it a governance challenge. Each cloud provider operates as a sovereign system with its own IAM model, policy language, role definitions, and permission granularity. When an organization deploys resources across three platforms, it effectively inherits three separate access control paradigms with no native inter-platform translation layer.

The Problem of IAM Heterogeneity

The IAM heterogeneity problem manifests at multiple levels. At the structural level, AWS uses explicit JSON policy documents that define allowed or denied actions on specific resources. Azure combines directory-based identity management through Azure Active Directory with scope-based role assignments across subscriptions, resource groups, and individual resources. GCP's IAM, in contrast, operates through a hierarchical model where inherited roles flow downward from organization to folder to project level.

They represent fundamentally different conceptual approaches to identity governance. A developer role in AWS is defined by a JSON policy listing permitted S3, EC2, and Lambda actions. The equivalent developer role in Azure is configured through RBAC assignments against a resource group. The GCP equivalent is a role binding at the project level. Mapping these three representations to each other requires not just technical translation but conceptual reconciliation and without automated tooling, this process is entirely manual.

Sitharaman, Karim, Gupta, and Tyagi (2025) identify that the heterogeneity of privilege models across cloud providers is "not merely a syntactical inconvenience but a structural barrier to scalable, secure access governance." Their research further establishes that while analytical frameworks for privilege analysis now exist, the field still lacks practical harmonization mechanisms capable of normalizing policies,

mapping equivalent roles, and generating actionable access recommendations across heterogeneous cloud providers a gap this study directly targets.

Beyond structural heterogeneity, multi-cloud environments are particularly vulnerable to privilege creep the gradual, often unnoticed accumulation of access rights as users change roles, take on temporary projects, or receive ad hoc permissions that are never revoked. In a multi-cloud context, a user may accumulate excessive permissions independently across three platforms over time, none of which appear excessive in isolation but which collectively represent a dramatically over-privileged access posture (Rouse, 2011; Tenable, 2023).

Excessive privileges expand the attack surface available to malicious actors who compromise user credentials. Overly permissive roles enable privilege escalation, lateral movement between cloud environments, accidental data exposure, and insider threat exploitation. In highly regulated sectors such as banking, healthcare, and government, these risks carry legal and financial consequences beyond immediate operational damage (CyberArk, 2024; Palo Alto Networks, 2024).

Despite the recognized severity of this problem, most organizations lack the tooling to detect and remediate it systematically. The absence of a centralized, cross-platform privilege visibility layer means that security teams must manually aggregate and analyze IAM data from each provider's separate dashboard an approach that is both labor-intensive and error-prone.

The academic and industry literature reveals several attempts to address multi-cloud IAM challenges. Sitharaman et al. (2025) propose a hypergraph-based analytical model capable of capturing multi-dimensional privilege dependencies across AWS and Azure, demonstrating measurably improved traversal and detection performance. However, this framework is purely analytical and stops short of privilege harmonization or actionable remediation.

Salah, Laborde, Benzekri, Kandi, and Ferreira (2025) explore Self-Sovereign Identity (SSI) as a solution for cross-cloud identity interoperability but focus on authentication rather than privilege management, with no practical harmonization mechanism. Mallesh (2025) proposes an AI-driven policy reconciliation system using transformer-based NLP models but presents only a conceptual architecture with no implemented system or evaluation results.

Gowda (2025) examines AWS and GCP identity governance using Infrastructure-as-Code and policy-as-code approaches, but omits Azure entirely and does not address individual user privilege harmonization. Faith (2025) surveys Privileged Access Management strategies across three platforms but delivers no implemented tooling, algorithms, or code.

Across this body of literature, a consistent pattern emerges: existing work gravitates toward either theoretical models or descriptive surveys, with very limited practical, implemented solutions. This research fills that gap.

Conceptual Framework of the Study

The conceptual framework of this study provides the theoretical and architectural map that guides the design of the Sentinel Harmonization Engine. It establishes the relationships between the core concepts multi-cloud computing, IAM, privilege management, access control models, and harmonization and defines how they interact within the proposed system.

The foundational premise of this study is that multi-cloud adoption, while operationally beneficial, inherently produces IAM fragmentation. When an organization deploys resources across AWS, Azure, and GCP simultaneously, it creates three independent identity silos with no native cross-platform communication. This fragmentation is the root cause of the privilege inconsistency problem that this system addresses.

The conceptual relationship is directional: multi-cloud adoption → IAM fragmentation → privilege inconsistency → security risk. Each stage of this chain amplifies the organizational security challenge, and each stage corresponds to a functional module in the Sentinel Harmonization Engine ingestion, normalization, detection, and remediation respectively.

This study is grounded in four established access control paradigms that collectively define the theoretical space within which the harmonization problem exists: Role-Based Access Control (RBAC) assigns permissions based on organizational roles rather than individual identities. It simplifies management at scale but can lack the flexibility needed for dynamic, multi-cloud environments. Azure's IAM model is primarily RBAC-based, making this the dominant paradigm for cross-platform role mapping in this system.

Attribute-Based Access Control (ABAC) evaluates access decisions based on contextual attributes including user identity, resource type, environmental conditions, and time. AWS's policy-based model incorporates ABAC elements through condition keys, enabling fine-grained, context-aware permission control.

Policy-Based Access Control, the primary model in AWS, defines permissions through structured JSON documents attached directly to users, groups, and roles. Its expressiveness enables fine-grained resource control but makes cross-platform standardization difficult.

The Zero Trust Security Model operates on the principle of "never trust, always verify." It demands continuous authentication, minimal standing privileges, and strict access validation at every interaction providing the security philosophy that underpins the least privilege principle enforced by this system.

The Sentinel Harmonization Engine does not implement any single model in isolation. Instead, it creates a meta-layer that translates across these models by normalizing their outputs into a unified HIGH / MEDIUM / LOW privilege scale, enabling comparison across paradigms that would otherwise be incommensurable.

The principle of least privilege (PoLP) is the normative goal that the entire harmonization process serves. It states that every user, process, or system component should operate with the minimum set of privileges necessary to perform its function. The Sentinel Harmonization Engine treats PoLP not merely as a theoretical objective but as a measurable outcome: the system's harmonization algorithm evaluates whether each user's privilege level across platforms is consistent, and flags users whose privilege levels in any environment exceed their established baseline as potential PoLP violations.

Privilege harmonization is defined in this study as the systematic process of:
(1) ingesting IAM data from heterogeneous cloud platforms;
(2) normalizing the data into a unified representational format;
(3) detecting inconsistencies between a user's privilege levels across platforms and
(4) generating remediation recommendations that bring all platforms into alignment with the least-privilege standard.

This four-stage process mirrors the modular architecture of the Sentinel Harmonization Engine. Each module in the system corresponds directly to one stage of the harmonization process, creating a conceptual alignment between the theoretical framework and the implemented artifact.

The Sentinel Harmonization Engine is implemented in Python and delivered through a Streamlit web dashboard. Its architecture comprises five integrated modules: The Data Ingestion Module accepts CSV-formatted IAM data files from AWS, Azure, and GCP simultaneously, detecting the platform format of each file automatically and parsing user privilege records accordingly.

The Normalization Engine maps each platform's native privilege representations onto a unified three-tier scale: HIGH (administrative or privileged access), MEDIUM (standard operational access), and LOW (read-only or minimal access). This normalization step is the technical core of the harmonization framework, resolving the semantic gap between AWS policy names, Azure RBAC roles, and GCP IAM bindings.

The Harmonization Algorithm compares each user's normalized privilege levels across all three platforms, identifies inconsistencies where privilege levels diverge, and calculates a risk score based on the severity and frequency of inconsistencies detected.

The Report Generation Module produces a Security Health Score (on a scale of 0 to 100), an executive summary of findings, detailed inconsistency reports with severity badges, and a privilege comparison matrix displaying each user's access level across all three platforms side by side.

The Interactive Dashboard presents all outputs through an authenticated Streamlit interface, providing real-time visual analysis, downloadable CSV and Excel reports, email alert simulation for critical findings, and a persistent scan history log.

Conceptual Variables

The study operates with the following key conceptual variables:

Independent Variable: IAM data ingested from AWS, Azure, and GCP (the multi-cloud privilege inputs).

Dependent Variable: Privilege inconsistency detection rate and Security Health Score (the outputs of the harmonization process).

Mediating Mechanism: The normalization engine and harmonization algorithm (the process by which inputs are transformed into actionable outputs).

This variable structure is consistent with Design Science Research methodology, which frames the system artifact as the mediating instrument between the identified problem (privilege inconsistency) and the desired outcome (harmonized, least-privilege access governance).

Nigeria, Cybersecurity Governance, and the Relevance of This Study

The Sentinel Harmonization Engine does not exist in a geopolitical vacuum. Its development at the Air Force Institute of Technology, Kaduna, reflects a deliberate engagement with the cybersecurity challenges facing Nigeria and the broader African technology ecosystem. Understanding the Nigerian policy environment and its relationship to multi-cloud security is essential for contextualizing the practical relevance of this work.

Nigeria's federal government has pursued an accelerating digital transformation agenda over the past decade, driven by policy frameworks including the National Digital Economy Policy and Strategy (NDEPS 2020–2030) and the National Cybersecurity Policy and Strategy (NCPS). These frameworks explicitly recognize cloud computing as a pillar of Nigeria's digital infrastructure and mandate security standards for government agencies adopting cloud services.

The adoption of cloud platforms by Nigerian government ministries, departments, and agencies (MDAs) has grown substantially, with institutions such as the National Identity Management Commission (NIMC), the Federal Inland Revenue Service (FIRS), and numerous state governments migrating significant workloads to AWS, Azure, and GCP. As these agencies expand their cloud footprints, the challenge of managing user privileges consistently across platforms becomes immediately relevant to national digital governance.

NITDA serves as Nigeria's primary regulatory body for information technology governance, including cloud security. The agency's Data Protection Regulation (NDPR) and its subsequent guidelines impose obligations on organizations processing Nigerian citizens' data to implement appropriate access controls directly implicating IAM management practices. Multi-cloud deployments that lack harmonized privilege governance create compliance exposure under NDPR,

particularly where excessive or inconsistent user privileges could enable unauthorized access to personal data.

The Sentinel Harmonization Engine directly supports NDPR compliance by providing a mechanism to detect and remediate privilege inconsistencies before they become compliance violations or data breach vectors.

Nigeria's financial services sector encompassing commercial banks, FinTech companies, payment processors, and insurance firms represents one of the most cloud-intensive industries in the country. The Central Bank of Nigeria (CBN) has issued Risk-Based Cybersecurity Framework guidelines that require financial institutions to maintain consistent access control policies and conduct regular privilege access reviews. Multi-cloud deployments complicate both requirements.

Financial institutions such as commercial banks operating on multiple cloud platforms face the exact privilege harmonization challenge this study addresses. An employee with administrative access in an AWS environment and read-only access in Azure presents a security inconsistency that the Sentinel Harmonization Engine would detect, classify, and recommend for remediation directly supporting CBN compliance obligations.

Nigerian universities, including AFIT itself, are increasingly deploying research workloads and administrative systems across cloud platforms. The challenge of managing student, faculty, and administrative staff access consistently across cloud environments is directly analogous to the enterprise use case this system addresses. By demonstrating a proof-of-concept implementation developed within a Nigerian academic institution, this study establishes that the tools and methodologies for multi-cloud privilege governance are accessible and applicable within the Nigerian context, not merely in resource-rich international organizations.

Nigeria faces a documented cybersecurity talent gap. The development of a functional, open-architecture privilege harmonization system by an undergraduate student at AFIT demonstrates the capacity of Nigerian academic institutions to produce practically relevant cybersecurity tooling. Beyond its immediate utility, the Sentinel Harmonization Engine serves as a training platform its modular architecture and documented codebase provide a foundation for future researchers, developers, and security practitioners to extend and build upon within the Nigerian higher education ecosystem.

The government's emphasis on Science, Technology, Engineering, and Mathematics (STEM) education and the operationalization of the National Cybersecurity Scholarship Program align directly with this kind of applied academic output. Research that produces functional security tooling rather than theoretical models alone represents the kind of indigenous capacity development that Nigeria's digital security ecosystem requires.

System Testing and Results

Testing of the Sentinel Harmonization Engine was conducted using simulated IAM datasets representing realistic multi-cloud privilege configurations across AWS, Azure, and GCP. The test environment included a dataset of twelve users with privilege assignments across all three platforms, incorporating known inconsistencies deliberately introduced to evaluate detection accuracy.

The system successfully ingested all three IAM data files, normalized privilege levels across platforms, and executed the harmonization algorithm without errors. Analysis identified seven distinct privilege inconsistencies across the twelve-user dataset. The Security Health Score computed by the system was 50 out of 100, correctly classifying the environment as a critical-risk privilege configuration given the density and severity of inconsistencies detected.

Findings included cases of privilege escalation (users with LOW privileges in GCP but HIGH privileges in AWS performing equivalent roles), privilege gaps (users with HIGH Azure access but no corresponding GCP permissions), and cross-platform misalignment in role definitions. The system generated severity-classified remediation recommendations for each inconsistency, distinguishing between CRITICAL, HIGH, and MEDIUM priority interventions.

The interactive dashboard presented these findings through a gauge-based Security Health Score visualization, a detailed inconsistency table with severity badges, a privilege comparison matrix enabling side-by-side three-platform analysis, and a before-and-after harmonization preview. Export functionality produced downloadable reports in both CSV and Excel formats. The email alert simulation successfully demonstrated alerting for critical-severity findings.

Conclusion

This study has demonstrated that automated, Python-based privilege harmonization across AWS, Azure, and GCP is technically achievable and practically valuable. The Sentinel Harmonization Engine fills a gap that existing literature focused largely on theoretical models and descriptive surveys has not addressed: a functional, implemented system that ingests real IAM data, detects privilege inconsistencies, computes risk scores, and generates actionable remediation guidance through an accessible web interface.

The system advances the state of multi-cloud IAM governance along several dimensions. It provides cross-platform visibility where none previously existed in a single tool. It translates heterogeneous IAM models AWS policy-based, Azure RBAC-based, and GCP hierarchy-based into a common representational framework. It operationalizes the principle of least privilege as a measurable, detectable, and remediable security property rather than an aspirational policy objective.

For Nigeria specifically, the system holds direct relevance across government agencies, financial institutions, higher education, and the growing FinTech sector, each of which faces multi-cloud IAM governance obligations under the NDPR, CBN guidelines, and NITDA frameworks. Its development at AFIT Kaduna demonstrates that world-class cybersecurity tooling can originate from Nigerian academic institutions a contribution to both national digital security capacity and the global academic discourse on cloud security governance.

Future work should extend the system to support real-time API integration with live cloud environments, expand the privilege classification schema to accommodate more granular role distinctions, incorporate machine learning for anomaly-based privilege drift detection, and introduce formal evaluation frameworks including confusion matrix analysis of detection accuracy across larger and more diverse datasets.

References

Ali, S., Khan, M., & Vasilakos, A. (2015). Security and privacy in cloud computing: Vision, trends, and challenges. Journal of Network and Computer Applications, 57, 1–10.

Ali, A., et al. (2025). Security and privacy in multi-cloud and hybrid cloud environments: A survey. Accepted manuscript.

Amajuoyi, C., Nwobodo, L., & Adegbola, A. (2024). Multi-cloud adoption and enterprise governance. International Journal of Cloud Computing Research.

Avirneni, D. (2025). Establishing workload identity for zero-trust CI/CD: From secrets to SPIFFE-based authentication. Technical Whitepaper.

Avadhani, P. (2025). Privilege creep in cloud environments: Detection and mitigation strategies. Journal of Cloud Security.

CyberArk. (2024). The 2024 identity security threat landscape report. CyberArk Inc.

Deochake, S., Murphy, R., & Gearheart, T. (2025). A multi-cloud framework for zero-trust workload authentication. IEEE Cloud Computing.

Ertl, B., Stevanovic, M., Hayrapetyan, A., Wegh, B., & Hardt, M. (2019). Identity harmonization for federated HPC, grid and cloud services. Future Generation Computer Systems, 94, 434–443.

Ezinwanneamaka, C., Kessie, J., Okaro, C., & Ezeife, C. (2025). Identity and access management in cloud storage: A comprehensive guide. Journal of Cloud Systems and Applications.

Faith, J. (2025). Multi-cloud privileged access management and just-in-time access. Cloud Security Review, 12(2).

Gowda, K. (2025). Governance at scale: Managing IAM and policy enforcement across AWS and GCP. International Journal of Software and Research in Emerging Technologies, 8(1).

Guo, T. (2025). Cross-cloud privilege inconsistency and its implications for enterprise security. ACM Computing Surveys.

Mallesh, N. (2025). Intelligent identity orchestration with AI-driven policy reconciliation for multi-cloud security. Journal of Artificial Intelligence in Cybersecurity.

Microsoft Corporation. (2025). Azure Active Directory and RBAC documentation. Microsoft. https://docs.microsoft.com/azure/

Nguyen, D. (2025). Cross-cloud IAM complexity and administrative overhead. Journal of Distributed Systems Security.

Palo Alto Networks. (2024). State of cloud-native security report 2024. Palo Alto Networks Inc.

Rouse, M. (2011). Privilege creep. TechTarget Security Encyclopedia.

Salah, S., Laborde, R., Benzekri, A., Kandi, M., & Ferreira, A. (2025). Identity management in cross-cloud environments: Towards self-sovereign identities. IEEE Transactions on Services Computing.

Sitharaman, S., Karim, R., Gupta, V., & Tyagi, A. K. (2025). Scalable privilege analysis for multi-cloud big data platforms: A hypergraph approach. IEEE Transactions on Cloud Computing, 13(1), 112–128.

Tenable. (2023). Cloud security report: Privilege management in multi-cloud environments. Tenable Inc.

Zhang, Q., Cheng, L., & Boutaba, R. (2010). Cloud computing: State-of-the-art and research challenges. Journal of Internet Services and Applications, 1(1), 7–18.

This article is derived from a final year project submitted to the Department of Cyber Security, Faculty of Computing, Air Force Institute of Technology (AFIT), Kaduna, Nigeria, in partial fulfilment of the requirements for the award of the Bachelor of Science degree in Cyber Security, June 2026.