INTERNATIONAL MULTI DISCIPLINARY JOURNAL FOR RESEARCH & DEVELOPMENT
eISSN 2394-6334 Impact factor: 7,854 Volume 13, issue 09 (2026)

Adaptive Access Control and Continuous Monitoring for Non-Human Identities in Cloud Environment

Amara Eze
Department of Computer Science, Nigerian Institute of Digital Technology, Nigeria

Abstract: The rapid adoption of cloud computing, distributed applications, automation, artificial intelligence, Internet of Things (IoT), and federated architectures has significantly increased the number of non-human identities (NHIs) operating within modern digital infrastructures. Service accounts, application identities, workload identities, machine credentials, automated agents, and API-based identities frequently perform privileged operations without direct human intervention. Consequently, conventional static access-control approaches are increasingly insufficient for environments in which identity behavior, workload context, and threat conditions change continuously. This research develops a conceptual framework for adaptive access control and continuous monitoring of non-human identities in cloud environments. The proposed approach integrates contextual authorization, behavioral monitoring, anomaly detection, federated intelligence, risk-based privilege adjustment, and automated response. The methodology synthesizes insights from research on federated learning, deep-learning-based intrusion detection, adversarial machine learning, privacy-preserving computation, and intelligent resource management. Particular emphasis is placed on the relationship between continuous behavioral observation and adaptive authorization decisions. The resulting framework treats NHI security as a dynamic lifecycle problem rather than a one-time authentication problem. The analysis indicates that adaptive controls can reduce the exposure associated with anomalous machine behavior while preserving operational flexibility, although false positives, model manipulation, privacy constraints, computational overhead, and incomplete behavioral context remain important limitations. The study contributes a research-oriented model for integrating identity governance with continuous security intelligence in cloud environments. networks.

Keywords: Non-Human Identities, Cloud IAM, Adaptive Access Control, Continuous Monitoring, Federated Learning, Behavioral Analytics, Anomaly Detection, Zero-Trust Security, Machine Identities

INTRODUCTION

The Cloud environments increasingly depend on identities that do not correspond to individual human users. Applications authenticate to databases, containers communicate with APIs, automation platforms execute infrastructure operations, and distributed workloads exchange credentials across dynamically changing environments. These identities may operate continuously and at machine speed, making their security characteristics fundamentally different from those of conventional human accounts. Ganapathy (2025) emphasizes the importance of automated governance and protection mechanisms for non-human identities in cloud IAM environments, particularly because the scale and operational complexity of machine identities make manual governance increasingly impractical.

The central security challenge is therefore not merely determining whether a non-human identity is authenticated, but determining whether its current behavior remains legitimate under changing environmental and operational conditions. A service account that legitimately accesses a storage service during normal operation may become suspicious if it suddenly initiates unusual access patterns, communicates with an unexpected service, or requests privileges inconsistent with its established workload behavior. Static authorization policies are poorly suited to such continuously changing circumstances because they generally

evaluate access according to predefined permissions rather than dynamically evolving risk.

Research on intelligent intrusion detection provides an important foundation for addressing this problem. Deep-learning-based anomaly detection has demonstrated the potential of behavioral models to identify deviations from expected activity, including in IoT networks (Alsoufi et al., 2024). Federated learning further provides an approach for collaboratively developing detection models without requiring all participating environments to centralize their raw data (Huang et al., 2024; Wen et al., 2023). These properties are particularly relevant to cloud environments containing multiple workloads, accounts, applications, and administrative domains.

The objective of this research is to develop a conceptual framework in which continuous monitoring informs adaptive access decisions for non-human identities. The research specifically investigates how behavioral intelligence can be integrated with identity authorization, how risk can dynamically influence privileges, and how federated and privacy-preserving techniques can support monitoring across distributed environments.

The scope is limited to the conceptual integration of the provided literature. The study does not claim empirical validation using a production cloud environment. Instead, it establishes a research framework that can subsequently be evaluated through simulation, controlled experiments, or real-world cloud deployments.

LITERATURE REVIEW

The literature provided for this study spans several technological domains that collectively establish the theoretical basis for adaptive NHI protection. Although the cited studies do not all directly investigate cloud non-human identities, their findings provide complementary foundations for behavioral security, distributed intelligence, anomaly detection, and adaptive decision-making.

Wen et al. (2023) examine federated learning through its challenges and applications, establishing the relevance of distributed model training where centralized collection of data may be undesirable or impractical. For cloud IAM, this principle is significant because NHI activity may be distributed across applications, environments, regions, and organizational boundaries. A centralized security-learning architecture can create data-consolidation challenges, whereas federated learning offers a conceptual mechanism for collaborative intelligence.

Huang et al. (2024) demonstrate the integration of hybrid deep-learning approaches with federated learning for intrusion detection. Their work is particularly relevant to the proposed framework because it connects distributed learning with security detection. Similarly, Priyadarshini (2024) examines IoT cyberattack detection through federated and split learning, reinforcing the potential of distributed intelligence for environments in which data processing and security analysis cannot always be concentrated in one location.

The literature also demonstrates the value of anomaly-based detection. Alsoufi et al. (2024) investigate deep learning for anomaly-based intrusion detection in IoT networks. This provides a theoretical basis for treating abnormal machine behavior as a security signal. In an NHI context, behavioral deviation can become an input to authorization decisions rather than merely an alert generated after an incident.

Adversarial considerations are equally important. He et al. (2023) examine adversarial machine learning for network intrusion detection systems, highlighting the security implications of relying on machine-learning models for defensive decisions. This introduces a critical issue for adaptive NHI governance: an attacker may attempt to manipulate behavioral patterns or exploit weaknesses in detection models to evade adaptive controls.

Ye et al. (2023) examine heterogeneous federated learning and its research challenges. This is relevant because cloud environments rarely exhibit homogeneous workloads. Different services may produce substantially different behavioral patterns. Consequently, an adaptive NHI framework must account for heterogeneity rather than assuming that one behavioral model applies equally to all identities.

Chen et al. (2024) address privacy-preserving computation in combination with federated learning. This is theoretically important because continuous monitoring requires extensive behavioral information. Without privacy safeguards, security monitoring itself may introduce excessive data exposure. The combination of distributed learning and privacy-preserving computation therefore offers a potential mechanism for reducing this tension.

Other provided studies contribute indirect but useful perspectives. Li et al. (2024) investigate intelligent edge resource scheduling using federated learning, demonstrating how distributed intelligence can support decision-making across resource-constrained environments. Alshingiti et al. and Khaleel et al. examine AI-based detection and identification approaches, while Mhmood et al. investigate deep-learning-based cyberattack detection. Salehi et al. (2023) provide a broader perspective on CNN and transfer learning, including advantages and challenges of learned representations. Oroojlooyjadid et al. demonstrate the application of deep reinforcement learning to optimization, providing a conceptual basis for considering adaptive decision policies.

The major research gap emerges from the lack of direct integration between behavioral security intelligence and NHI authorization governance within the supplied literature. Existing works provide pieces of the solution—federated learning, anomaly detection, privacy preservation, adversarial analysis, and intelligent decision-making—but a unified framework connecting these capabilities to continuously changing NHI privileges remains insufficiently developed. Ganapathy (2025) directly reinforces the importance of automated governance and protection of NHIs in cloud IAM, providing the identity-governance foundation upon which the proposed framework is positioned.

METHODOLOGY

3.1 Research Design

This study adopts a conceptual research methodology based on structured synthesis of the provided literature. The methodological objective is to transform findings from distributed learning, anomaly detection, adversarial machine learning, privacy-preserving computation, and intelligent decision-making into a unified adaptive-access architecture for NHIs.

The proposed methodology follows five logical stages: identity profiling, contextual evaluation, behavioral monitoring, risk computation, and adaptive enforcement. These stages operate continuously rather than sequentially as a one-time authentication process.

3.2 Non-Human Identity Profiling

The first stage establishes a behavioral and authorization profile for each NHI. Rather than representing an identity solely through its assigned permissions, the profile incorporates its expected operational context.

A conceptual profile can be represented as:

Pi={Ii,Ri,Wi,Si,Bi,Ti}P_i = \{I_i, R_i, W_i, S_i, B_i, T_i\}

where IiI_i represents identity attributes, RiR_i represents authorized resources, WiW_i represents workload context, SiS_i represents service relationships, BiB_i represents behavioral characteristics, and TiT_i represents temporal operating patterns.

For example, an automated deployment identity may normally access a container registry, infrastructure-management API, and deployment platform during predefined operational windows. A request outside this behavioral envelope should not automatically be classified as malicious, but it should contribute to a changing risk score.

This approach extends conventional IAM governance toward the automated governance model discussed by

Ganapathy (2025).

3.3 Context-Aware Access Evaluation

The second stage evaluates each access request against both authorization rules and contextual conditions. The conceptual decision function is:

$D=f(Pi,Ct,Bt,Rt)D = f(P_i,C_t,B_t,R_t)$

where PiP_i is the identity profile, CtC_t is current contextual information, BtB_t represents observed behavior, and RtR_t represents the current risk state.

The framework therefore distinguishes between permission eligibility and access suitability. An identity may possess permission to access a resource but still receive a restricted decision if its current behavior substantially deviates from its expected profile.

This distinction is essential for adaptive access control because permissions represent intended capability, whereas behavioral evidence provides information about current security conditions.

3.4 Continuous Behavioral Monitoring

Continuous monitoring forms the central analytical component of the proposed architecture. Events such as authentication attempts, API calls, resource access, geographic or network context, privilege changes, service interactions, and access frequency can be transformed into behavioral features.

Anomaly detection models can then compare current behavior with historical or learned patterns. The relevance of this approach is supported by anomaly-based deep-learning research such as Alsoufi et al. (2024), while the use of federated learning for security detection is supported by Huang et al. (2024) and Priyadarshini (2024).

The framework does not assume that every anomaly represents an attack. Instead, anomalies act as risk signals. This distinction reduces the possibility of excessively aggressive access blocking.

For example, if an NHI normally generates 100 API requests per hour but suddenly generates 10,000 requests, the monitoring layer can increase its risk score. The access-control layer may then reduce privileges, require additional verification mechanisms where applicable, or isolate selected operations.

3.5 Federated Security Intelligence

Cloud environments may contain multiple security domains whose behavioral data should not necessarily be centralized. Federated learning provides a conceptual mechanism for training models across distributed participants while reducing direct exchange of raw datasets (Wen et al., 2023).

A federated NHI monitoring architecture can therefore allow different cloud environments to contribute model updates rather than complete activity datasets. Huang et al. (2024) and Priyadarshini (2024) demonstrate the relevance of federated approaches to distributed security detection.

However, heterogeneous learning environments introduce additional complexity. Ye et al. (2023) identify challenges associated with heterogeneous federated learning. NHI workloads similarly exhibit substantial behavioral diversity. A database service identity cannot necessarily be modeled using the same behavioral assumptions as an automated CI/CD identity.

Consequently, the proposed framework recommends identity-class-specific or workload-aware behavioral models rather than one universal model.

3.6 Privacy-Preserving Monitoring

Continuous monitoring creates an inherent privacy and governance tension: better detection generally requires richer behavioral information. Chen et al. (2024) provide a theoretical basis for combining federated learning with privacy-preserving computation to reduce this tension.

The proposed architecture therefore separates raw telemetry from learned security intelligence wherever practical. Rather than transmitting complete activity histories, participating environments can contribute protected model information.

The limitation is that privacy mechanisms may increase computational complexity and can potentially reduce the fidelity of security analysis. Consequently, privacy should be treated as a design constraint rather than an independent feature.

3.7 Risk-Based Adaptive Enforcement

The final methodological stage converts behavioral intelligence into authorization decisions. A conceptual risk function can be expressed as:

Riskt=w1At+w2Vt+w3Dt+w4CtRisk_t = w_1A_t+w_2V_t+w_3D_t+w_4C_t

where AtA_t represents anomaly severity, VtV_t represents privilege sensitivity, DtD_t represents behavioral deviation, and CtC_t represents contextual risk.

The resulting score can map to graduated responses:

Riskt→{Allowlow risk Restrictmoderate risk Challenge/ Verifyhigh risk Suspend/ Isolatecritical risk Risk_t \rightarrow \begin{cases} \text{Allow} & \text{low risk}\\ \text{Restrict} & \text{moderate risk}\\ \text{Challenge/Verify} & \text{high risk}\\ \text{Suspend/Isolate} & \text{critical risk} \end{cases}

This graduated approach is preferable to binary authorization because cloud workloads frequently experience legitimate behavioral variation.

Deep reinforcement learning research such as Oroojlooyjadid et al. provides a conceptual foundation for considering intelligent optimization of decision policies, although direct application to NHI access control would require separate empirical validation.

RESULTS

The conceptual analysis produces several important findings regarding adaptive NHI protection.

First, continuous monitoring should be treated as an authorization input rather than merely a detection mechanism. Conventional security architectures may generate an alert after abnormal behavior has occurred, whereas the proposed model incorporates behavioral evidence directly into subsequent access decisions. This transforms monitoring from a passive observability function into an active governance mechanism.

Second, the analysis indicates that adaptive access control is most effective when privileges are dynamically associated with contextual risk. A permanently authorized service account represents a substantial security exposure if its credentials are compromised. Conversely, dynamically restricting access when behavioral risk increases can reduce the potential impact of credential misuse. This direction is consistent with the automated NHI governance perspective presented by Ganapathy (2025).

Third, federated learning provides a potentially valuable architecture for distributed NHI intelligence. Research by Wen et al. (2023), Huang et al. (2024), and Priyadarshini (2024) demonstrates the applicability of distributed learning to security-oriented scenarios. However, the findings also indicate that heterogeneous workloads require specialized modeling strategies rather than uniform behavioral assumptions.

Fourth, anomaly detection should not operate as an isolated binary classifier. The literature on anomaly

detection and adversarial machine learning indicates that machine-learning-based detection introduces both opportunities and vulnerabilities (Alsoufi et al., 2024; He et al., 2023). Consequently, the proposed architecture uses anomaly scores as one component of a broader risk calculation.

Fifth, privacy-preserving techniques are important because continuous monitoring creates extensive telemetry. The combination of federated learning and privacy-preserving computation identified by Chen et al. (2024) offers a potential pathway for balancing detection effectiveness with data-minimization objectives.

Finally, the findings indicate that adaptive governance introduces operational trade-offs. Aggressive privilege reduction may improve security but can interrupt legitimate automated workloads. Conversely, excessively permissive thresholds can reduce the value of continuous monitoring. Therefore, effective implementation requires calibrated risk thresholds, identity-specific baselines, explainable decisions, and mechanisms for controlled recovery.

DISCUSSION

The primary theoretical contribution of the proposed framework is its treatment of NHI security as a continuous control problem rather than a static authentication problem. Human-centered IAM frequently emphasizes authentication and predefined authorization, whereas machine identities can operate continuously and at significantly greater scale. Their behavior therefore becomes an important security attribute.

The literature supports the technical components of this perspective but also exposes significant limitations. Federated learning can distribute intelligence, yet heterogeneous environments create difficulties in developing generalized models (Ye et al., 2023). Anomaly detection can identify behavioral deviations, yet abnormality does not necessarily imply malicious activity. Deep-learning approaches can improve detection capabilities, but adversarial machine learning demonstrates that defensive models themselves can become attack surfaces (He et al., 2023).

This creates an important trade-off between security sensitivity and operational stability. If every unusual behavior triggers immediate suspension, legitimate workload changes may cause service disruptions. If the framework waits for high-confidence malicious evidence, attackers may retain access long enough to cause damage. Adaptive controls should therefore operate through graduated interventions rather than a single allow-or-deny threshold.

Another theoretical implication concerns privacy. Continuous NHI monitoring can generate highly detailed behavioral records. Although Chen et al. (2024) demonstrate the relevance of privacy-preserving computation in distributed learning, privacy mechanisms may introduce additional computational requirements. The system must therefore balance privacy, detection accuracy, latency, and operational cost.

The framework also has practical implications for cloud governance. Ganapathy (2025) positions automated governance as an important component of NHI protection. The present research extends this perspective by connecting governance with continuous behavioral evidence. Instead of asking only whether an identity has an approved permission, the system continually evaluates whether the identity's current behavior remains consistent with its expected operational purpose.

The framework has several limitations. First, it is conceptual and has not been validated against production cloud telemetry. Second, the supplied literature primarily addresses adjacent domains rather than providing extensive empirical evidence specifically focused on cloud NHIs. Third, the proposed risk function is illustrative rather than experimentally calibrated. Fourth, false positives and model drift may significantly affect operational effectiveness. Finally, adversaries capable of manipulating training data or behavioral patterns may undermine adaptive models.

Future empirical research should therefore construct representative NHI datasets, evaluate different behavioral models, measure detection accuracy and false-positive rates, and examine the latency and resource overhead of adaptive authorization. Such experiments would provide the evidence required to determine whether the

proposed conceptual architecture delivers measurable improvements over static IAM approaches.

CONCLUSION

This research presented a conceptual framework for adaptive access control and continuous monitoring of non-human identities in cloud environments. The analysis demonstrates that securing machine identities requires more than static permissions and periodic reviews. Because NHIs operate continuously across dynamic cloud workloads, their current behavioral context should contribute directly to access-control decisions.

The proposed framework integrates identity profiling, contextual authorization, continuous behavioral monitoring, anomaly detection, federated learning, privacy-preserving intelligence, and risk-based adaptive enforcement. The literature indicates that federated learning can support distributed security intelligence, anomaly detection can identify behavioral deviations, and privacy-preserving computation can reduce some of the risks associated with centralized telemetry.

The principal contribution is the conceptual integration of these capabilities into an identity-governance model in which monitoring and authorization form a continuous feedback loop. Low-risk behavior can preserve normal operational privileges, while increasing behavioral risk can trigger graduated restrictions or isolation. This approach provides a more dynamic alternative to purely static authorization models.

Nevertheless, adaptive access control should not be considered a complete solution without empirical validation. Model drift, adversarial manipulation, false positives, privacy constraints, workload heterogeneity, and operational overhead remain important challenges. Future research should validate the framework using realistic cloud workloads and develop measurable policies for balancing security, availability, privacy, and computational efficiency.

REFERENCES

1. Z. Alshingiti et al., “A Deep Learning-based Phishing Detection System Using CNN, LSTM, and LSTM-CNN,” Electronics, vol. 12, no. 1.

2. M. A. Alsoufi et al., “Anomaly-Based Intrusion Detection Model Using Deep Learning for IoT Networks,” Computer Modeling in Engineering & Sciences, vol. 141, no. 1, pp. 823–845, 2024.

3. F. Calero et al., “A Review of Modeling and Applications of Energy Storage Systems in Power Grids,” Proceedings of the IEEE, vol. 111, no. 7, pp. 806–831, 2022.

4. J. Chen et al., “When Federated Learning Meets Privacy-preserving Computation,” ACM Computing Surveys, vol. 56, no. 12, pp. 1–36, 2024.

5. K. He et al., “Adversarial Machine Learning for Network Intrusion Detection Systems: A Comprehensive Survey,” IEEE Communications Surveys & Tutorials, vol. 25, no. 1, pp. 538–566, 2023.

6. J. Huang et al., “Improved Intrusion Detection Based on Hybrid Deep Learning Models and Federated Learning,” Sensors, vol. 24, no. 12, Art. no. 4002, 2024.

7. M. Khaleel et al., “Artificial Intelligent Techniques for Identifying the Cause of Disturbances in the Power Grid,” Brilliance: Research of Artificial Intelligence, vol. 3, no. 1, pp. 19–31.

8. H. Li et al., “Optimizing Intelligent Edge Computing Resource Scheduling Based on Federated Learning,” Journal of Knowledge Learning and Science Technology, vol. 3, no. 3, pp. 235–260.

9. M. S. Mastoi et al., “Large-scale Wind Power Grid Integration Challenges and Their Solution: A

Detailed Review," Environmental Science and Pollution Research, vol. 30, no. 47, pp. 103424–103462, 2023.

10. A. Mhmood et al., "Detection of Cyber-attacks on Smart Grids Using Improved VGG19 Deep Neural Network Architecture and Aquila Optimizer Algorithm," Signal, Image and Video Processing, vol. 18, no. 2, pp. 1477–1491, 2024.

11. Orooj looyjadid et al., "A Deep Q-network for the Beer Game: Deep Reinforcement Learning for Inventory Optimization," Manufacturing & Service Operations Management, vol. 24, no. 1, pp. 285–304.

12. Priyadarshini, "Anomaly Detection of IoT Cyberattacks in Smart Cities Using Federated Learning and Split Learning," Big Data and Cognitive Computing, vol. 8, no. 3, Art. no. 21, 2024.

13. A. W. Salehi et al., "A Study of CNN and Transfer Learning in Medical Imaging: Advantages, Challenges, Future Scope," Sustainability, vol. 15, no. 7, Art. no. 5930, 2023.

14. J. Wen et al., "A Survey on Federated Learning: Challenges and Applications," International Journal of Machine Learning and Cybernetics, vol. 14, no. 2, pp. 513–535, 2023.

15. M. Ye et al., "Heterogeneous Federated Learning: State-of-the-art and Research Challenges," ACM Computing Surveys, vol. 56, no. 3, pp. 1–44.

16. Ganapathy, S. K. (2025). Automated Governance and Protection of Non-Human Identities in Cloud IAM. Frontiers in Emerging Computer Science and Information Technology, 2(02), 08–20. Retrieved from https://irjernet.com/index.php/fecsit/article/view/cloud-iam-non-human-identities