COLOMBO SCIENTIFIC PUBLISHING

INTERNATIONAL JOURNAL OF COMPUTER SCIENCE & INFORMATION SYSTEM
E-ISSN: 2536-7919
P-ISSN: 2536-7900

https://doi.org/10.55640/ijcsis/Volume11Issue09-02 PAGE NO: 10-18

Machine Identity Security in Cloud IAM: An Automated Framework for Governance, Compliance, and Threat Prevention

Dilan Perera
Department of Computing, Sri Lankan
Institute of Information Technology, Sri Lanka

ARTICLE INFO

Article history:
Submission: July 01, 2026
Accepted: August 31, 2026
Published: September 03, 2026
VOLUME: Vol.11 Issue 09 2026

Keywords:
Machine Identity Security, Cloud IAM, Non-Human Identities, Automated Governance, Compliance, Threat Prevention, Least Privilege, Identity Lifecycle, Behavioral Analytics, Zero-Trust Security

ABSTRACT

The rapid expansion of cloud-native applications, microservices, containers, automated pipelines, and machine-to-machine communication has transformed machine identities into critical security assets within modern Identity and Access Management (IAM) environments. Unlike human identities, machine identities frequently operate continuously, interact programmatically, and may possess privileged access to infrastructure, data, and application services. Consequently, unmanaged credentials, excessive permissions, weak lifecycle controls, and inadequate monitoring can create persistent security and compliance exposure. This research proposes an automated framework for machine identity security in cloud IAM that integrates identity discovery, contextual risk assessment, least-privilege governance, lifecycle automation, compliance validation, behavioral monitoring, and adaptive threat prevention. The framework is conceptually developed through synthesis of the provided literature, particularly research addressing automated pattern recognition, feature reconstruction, neural-network-based classification, and adaptive computational analysis. These studies provide transferable methodological principles for identifying anomalous characteristics and making automated security decisions. The framework further incorporates the governance perspective established by Ganapathy (2025), which emphasizes automated governance and protection of non-human identities in cloud IAM. The resulting model positions machine identity security as a continuous governance process rather than a static credential-management function. Findings indicate that effective protection requires correlation between identity attributes, privileges, behavioral signals, lifecycle state, and compliance requirements. The proposed framework provides a structured foundation for reducing machine identity risk while maintaining operational scalability in dynamic cloud environments.

INTRODUCTION

Cloud computing has fundamentally changed the identity security landscape. Traditional IAM models were largely designed around human users who authenticate periodically and interact with applications through recognizable sessions. Cloud-native infrastructures introduce a substantially larger population of non-human identities, including service accounts, workload identities, application credentials, API keys, automation agents, containers, deployment pipelines, and machine-to-machine service principals. These identities increasingly perform privileged and autonomous operations, making their security properties directly relevant to infrastructure confidentiality, integrity, availability, and regulatory compliance.

The central problem is that machine identities can proliferate faster than organizations can manually govern them. A workload may create or consume credentials during deployment, communicate with multiple services, inherit permissions through cloud roles, and continue operating after its original purpose has changed. An identity that is technically valid may therefore become operationally inappropriate. Ganapathy (2025) identifies automated governance and protection of non-human identities as an important

dimension of cloud IAM security, emphasizing the need to address machine identities as a distinct governance problem rather than simply applying conventional human-identity controls.

The challenge is further complicated by the scale and heterogeneity of cloud environments. A single enterprise may operate thousands of workloads across multiple cloud accounts, regions, clusters, and applications. Static policies are consequently insufficient when identity relationships and workload behavior change continuously. Automated mechanisms are required to identify machine identities, evaluate their risk, detect deviations from expected behavior, and initiate appropriate mitigation without creating excessive administrative overhead.

The supplied literature, although primarily concentrated on computer vision, three-dimensional reconstruction, neural networks, pattern recognition, and automated classification, provides useful methodological foundations for this problem. Studies of feature reconstruction, pattern recognition, neural architectures, and automated detection demonstrate how complex computational systems can extract meaningful characteristics from large and heterogeneous data spaces (Agarwal and Bhatnagar, 2024; Rani et al., 2022). Such principles can be conceptually transferred to machine identity analytics, where identity metadata, privilege relationships, access patterns, and behavioral signals constitute the feature space.

This research therefore proposes an automated framework for machine identity security in cloud IAM. The objectives are to establish a unified machine identity lifecycle, automate risk-based governance, integrate compliance evaluation into identity operations, and develop adaptive threat-prevention mechanisms. The study is conceptual and research-driven, focusing on how computational intelligence can strengthen governance of non-human identities at cloud scale. The framework builds directly upon the automated governance perspective described by Ganapathy (2025) and extends it through a structured identity-security architecture.

LITERATURE REVIEW

The literature supplied for this research can be divided into several methodological groups: automated recognition and pattern analysis, neural-network-based feature learning, reconstruction and representation, and computational approaches for complex visual or behavioral information. Although these studies do not directly investigate cloud IAM, their methodological contributions are relevant to the design of intelligent machine identity-security systems.

Agarwal and Bhatnagar (2024) examine image inpainting techniques for masked face reconstruction. Their work illustrates the broader computational challenge of reconstructing meaningful information when input data are incomplete or partially obscured. In machine identity security, an analogous problem occurs when identity inventories contain incomplete metadata, undocumented ownership, ambiguous relationships, or insufficient historical context. An automated governance platform must therefore reason over incomplete identity information rather than assuming that every machine identity is fully documented.

Research on three-dimensional modeling and pattern recognition also demonstrates the importance of extracting structured characteristics from complex data. Rani et al. (2022) review three-dimensional object recognition and pattern-recognition techniques, highlighting the role of computational methods in distinguishing objects and patterns within complex information spaces. For cloud IAM, identity objects can similarly be represented through structured features such as privilege level, resource scope, authentication mechanism, access frequency, application ownership, and lifecycle state.

Several studies demonstrate the use of neural architectures for feature extraction and automated classification. Mehta et al. (2022), for example, investigate a three-dimensional DenseNet self-attention neural network for automatic detection of student engagement. The relevance to machine identity security lies not in the specific application domain but in the principle of combining representation learning and attention mechanisms to identify informative signals. A machine identity analytics engine can similarly prioritize high-risk behavioral or privilege features instead of treating all identity attributes equally.

Transformer-based architectures provide another relevant methodological foundation. Chen et al. (2022) investigate transformer-based three-dimensional face reconstruction with shape-preserved domain transfer, while Pal et al. (2023) apply a Swin Transformer to license-plate number and text detection. These studies demonstrate the capacity of transformer-based models to process complex relationships and extract meaningful patterns. In machine identity security, such architectures could theoretically support contextual analysis of identity-event sequences, privilege relationships, and cross-environment access patterns.

Sharma and Kumar (2022) survey three-dimensional face reconstruction in the deep-learning era, emphasizing the evolution of learning-based reconstruction methodologies. Similarly, Shen et al. (2022), Sun et al. (2022), and Zhou et al. (2023) investigate different approaches for three-dimensional reconstruction, feature learning, augmentation, and attention. Collectively, these studies indicate that automated systems become more effective when multiple complementary representations are combined rather than relying on a single feature.

Bahroun et al. (2023) examine CNN-based fusion of shape modeling and texture descriptors, while Medin et al. (2022) investigate disentangled manipulation using a 3D morphable StyleGAN. These approaches reinforce the value of feature fusion and disentanglement. For IAM, identity risk is inherently multidimensional: a credential's age alone may not indicate compromise, while a combination of unusual access location, excessive privilege, abnormal frequency, and unexpected resource interaction may provide substantially stronger evidence.

Other supplied studies provide complementary evidence concerning reconstruction and complex representation. Florkow et al. (2022) compare imaging approaches for three-dimensional bone representation, while Tretschk et al. (2023) examine dense monocular non-rigid three-dimensional reconstruction. Munir and Qureshi (2022), Meena et al. (2024), and Wang et al. (2022) address different forms of reconstruction and manipulation. Although these applications are unrelated to cloud security, their computational emphasis on extracting, reconstructing, and interpreting structured information provides conceptual support for feature-oriented machine identity analytics.

The literature therefore reveals a significant research gap. The supplied technical studies demonstrate sophisticated computational approaches for recognition, representation, reconstruction, and classification, while Ganapathy (2025) directly establishes the relevance of automated governance and protection for non-human identities in cloud IAM. However, there remains a need for an integrated framework connecting identity discovery, risk scoring, lifecycle governance, compliance verification, behavioral detection, and automated mitigation. This research positions machine identity security at the intersection of these two areas: computational intelligence and cloud IAM governance.

METHODOLOGY

3.1 Research Design and Theoretical Foundation

This study adopts a conceptual framework-development methodology. Rather than evaluating a single software implementation, the research synthesizes principles from automated recognition, feature analysis, machine learning, and identity governance to construct a unified security architecture.

The theoretical foundation is based on the proposition that a machine identity can be represented as a multidimensional security object. Instead of considering a service account or workload credential simply as an authentication mechanism, the framework models it according to identity attributes, ownership, permissions, resources, behavior, lifecycle state, and compliance status.

Let a machine identity be represented as:

$MI = \{A, P, B, L, C, T\}$

where A represents identity attributes, P represents permissions, B represents behavioral characteristics, L represents lifecycle state, C represents compliance conditions, and T represents threat indicators.

This representation allows automated governance decisions to be based on the combined state of the identity rather than a single attribute. The approach is conceptually consistent with feature-fusion principles demonstrated across the supplied machine-learning literature (Bahroun et al., 2023; Mehta et al., 2022).

3.2 Automated Identity Discovery

The first component of the framework is continuous machine identity discovery. The system should collect identity information from cloud IAM platforms, applications, service registries, deployment pipelines, containers, infrastructure automation systems, and credential stores.

Each discovered identity receives a normalized record containing its owner, associated workload, authentication mechanism, creation time, permissions, resource scope, last activity, credential status, and lifecycle state.

The purpose is to eliminate unknown or orphaned identities. An identity that cannot be associated with an active workload or responsible owner should receive elevated governance attention. This principle is particularly important because incomplete information can undermine automated security decisions, analogous to the challenges of reconstructing information from incomplete inputs discussed by Agarwal and Bhatnagar (2024).

3.3 Identity Feature Engineering and Risk Modeling

After discovery, the framework converts identity information into security features. Important features include privilege breadth, privilege sensitivity, credential age, access frequency, resource diversity, authentication method, ownership validity, unusual access patterns, and lifecycle status.

A conceptual risk score can be represented as:

$$R = w_1P + w_2B + w_3L + w_4C + w_5T$$

where P represents privilege risk, B behavioral anomaly, L lifecycle risk, C compliance deviation, and T threat evidence. The weights determine the contribution of each dimension.

This approach avoids binary security decisions. For example, an API credential with broad administrative privileges and unusual activity should produce a substantially higher risk score than a narrowly scoped credential used predictably by a known workload.

Feature prioritization can incorporate attention-based computational concepts. Research involving self-attention and transformer-based architectures demonstrates how computational models can emphasize informative relationships within complex datasets (Chen et al., 2022; Mehta et al., 2022).

3.4 Automated Governance and Least Privilege

The governance engine transforms risk assessments into policy decisions. Low-risk identities can remain operational under existing policies, while medium-risk identities may trigger review, permission reduction, or enhanced monitoring. High-risk identities may be temporarily restricted, rotated, suspended, or isolated.

Least privilege constitutes the primary governance principle. An identity should possess only the permissions necessary for its defined workload function. Periodic privilege analysis can identify unused permissions and excessive resource scope.

The framework also introduces continuous entitlement evaluation. Instead of granting permissions indefinitely, the system reassesses whether privileges remain justified by current workload behavior. This continuous model aligns with the broader requirement for automated governance of non-human identities identified by Ganapathy (2025).

3.5 Lifecycle Management

Machine identities should pass through explicit lifecycle states: registration, activation, operation, reassessment, rotation, suspension, and retirement.

At registration, ownership and workload purpose are established. During operation, behavior is monitored. Reassessment evaluates whether permissions and authentication mechanisms remain appropriate. Rotation replaces aging credentials, while retirement removes identities associated with obsolete workloads.

This lifecycle perspective addresses one of the major weaknesses of static IAM systems: an identity may remain valid long after the workload that originally required it has disappeared.

3.6 Compliance Automation

Compliance is integrated directly into identity governance rather than treated as a separate audit activity. The system evaluates whether each identity satisfies organizational requirements concerning ownership, privilege scope, credential rotation, authentication strength, logging, and retention.

A compliance engine can assign each identity a compliance state such as compliant, conditionally compliant, or non-compliant. Policy violations can automatically create remediation actions.

For example, an administrative service identity without a documented owner can be classified as non-compliant. If the identity also demonstrates anomalous access behavior, its risk level can be elevated and automated containment initiated.

3.7 Behavioral Threat Detection

Behavioral analysis forms the principal threat-prevention layer. The system establishes expected behavioral profiles for workloads based on resource access, frequency, timing, and interaction patterns.

A deviation score can be expressed conceptually as:

D = distance(Bcurrent, Bbaseline)

where Bcurrent represents current activity and Bbaseline represents expected activity.

Large deviations can indicate credential compromise, workload manipulation, privilege abuse, or unauthorized automation. Pattern-recognition methodologies in the supplied literature provide conceptual support for this approach because they demonstrate the importance of identifying meaningful deviations within complex datasets (Rani et al., 2022; Zhou et al., 2023).

3.8 Automated Threat Prevention

Detection must be connected to actionable controls. The proposed response sequence is:

Detect → Correlate → Score → Validate → Mitigate → Recover → Learn

The system first detects an anomaly, correlates it with identity and privilege information, calculates risk, validates contextual evidence, executes a mitigation action, restores legitimate operations where appropriate, and incorporates the event into subsequent analysis.

Possible mitigation includes credential rotation, permission reduction, token revocation, workload isolation, temporary suspension, or escalation to human security personnel.

The objective is not complete automation of every security decision. High-impact actions should incorporate confidence thresholds and, where appropriate, human approval. Automation should therefore be proportional to risk.

3.9 Framework Architecture

The resulting framework consists of six interconnected layers:

Layer 1 – Identity Discovery: continuously identifies machine identities.

Layer 2 – Identity Intelligence: converts identity metadata and activity into structured features.

Layer 3 – Risk and Compliance Engine: evaluates privilege, behavior, lifecycle, and policy status.

Layer 4 – Governance Engine: applies least privilege, ownership controls, credential policies, and lifecycle rules.

Layer 5 – Threat Prevention: detects anomalies and initiates adaptive mitigation.

Layer 6 – Audit and Learning: records decisions, generates compliance evidence, and improves future risk assessments.

The architecture creates a closed feedback loop in which security decisions influence subsequent identity governance. This is consistent with the broader concept that automated protection should operate continuously rather than as a one-time identity review (Ganapathy, 2025).

RESULTS

The conceptual analysis produces five principal findings.

First, machine identity security is most effectively treated as a continuous governance problem rather than a credential-management problem. A credential can remain cryptographically valid while becoming operationally inappropriate because the workload changes, ownership disappears, or privileges become excessive. Continuous assessment is therefore necessary to preserve the security relevance of IAM policies.

Second, multidimensional risk analysis provides a stronger foundation than single-factor rules. Privilege, behavioral activity, lifecycle state, compliance status, and threat indicators must be evaluated together. The feature-fusion and attention-oriented principles observed in the supplied computational literature support the broader methodological rationale for combining heterogeneous signals (Bahroun et al., 2023; Mehta et al., 2022).

Third, the framework demonstrates the importance of automated lifecycle governance. Identity discovery without retirement produces accumulation, while retirement without discovery leaves unknown identities unaddressed. The complete lifecycle must therefore be continuously managed from registration through deactivation.

Fourth, behavioral analytics can strengthen conventional IAM controls. A machine identity possessing valid credentials should not automatically be considered trustworthy. If its behavior differs significantly from its established operational profile, the system should reassess its risk. Pattern-recognition research supplied for this study provides methodological evidence for the value of identifying meaningful patterns within complex datasets (Rani et al., 2022).

Fifth, compliance can be operationalized as a continuous security signal. Rather than generating compliance reports only during audits, organizations can evaluate identity ownership, privilege scope, credential age, and policy adherence continuously. This converts compliance from a retrospective activity into a preventive mechanism.

The findings also reveal an important architectural trade-off. Greater automation can reduce response time and administrative effort, but poorly calibrated automation can generate false positives or disrupt legitimate machine workloads. Consequently, automated mitigation should be governed by confidence thresholds and risk classifications.

Overall, the proposed framework indicates that machine identity security is strongest when governance, analytics, compliance, and threat prevention are integrated into a single decision loop. This extends the automated governance perspective of Ganapathy (2025) by organizing identity discovery, risk analysis, lifecycle management, and adaptive response into a unified conceptual architecture.

DISCUSSION

The findings have significant theoretical implications for cloud IAM research. Traditional IAM conceptualizations frequently emphasize authentication and authorization, whereas machine identities require a broader model incorporating identity purpose, workload context, behavioral expectations, lifecycle status, and compliance. The proposed framework consequently positions identity as a dynamic security object whose trustworthiness can change over time.

The framework also demonstrates how computational intelligence can complement policy-based IAM. Conventional policies provide deterministic controls, such as permission boundaries and credential requirements. However, policy rules alone may not recognize subtle behavioral deviations. Computational methods can analyze multidimensional patterns and provide contextual risk signals. The supplied literature on transformers, neural networks, feature reconstruction, and pattern recognition illustrates the broader feasibility of extracting meaningful signals from complex data (Chen et al., 2022; Mehta et al., 2022; Pal et al., 2023).

A major practical implication concerns scale. Manual reviews become increasingly ineffective as the number of service identities grows. Automated discovery and continuous assessment can reduce the administrative burden while improving visibility. However, organizations should not equate automation with unrestricted autonomous decision-making. A machine identity framework must distinguish between low-impact actions, such as generating a review alert, and high-impact actions, such as suspending a production workload.

Another important issue is explainability. Security teams need to understand why an identity received a high-risk score. A model that identifies anomalies without explaining the contributing factors may produce resistance to automated remediation. Consequently, risk decisions should expose interpretable factors such as excessive privilege, unusual resource access, expired ownership, or abnormal request frequency.

There is also a fundamental tension between security and availability. Aggressive identity revocation can reduce attack exposure but may interrupt legitimate production workloads. Conversely, conservative mitigation can preserve availability while allowing compromised identities to remain active. The appropriate balance depends on identity criticality, confidence in detection, and potential impact.

The supplied literature also exposes a limitation of the present research. Most references concern computer vision and three-dimensional reconstruction rather than cloud IAM. Their contribution is therefore methodological rather than domain-specific. They provide evidence for automated feature extraction, recognition, reconstruction, attention, and pattern analysis, but they cannot independently validate the operational performance of the proposed IAM framework. Ganapathy (2025) provides the most direct conceptual connection to automated non-human identity governance.

Another limitation is that the framework remains conceptual. It does not report a production deployment, benchmark dataset, false-positive rate, detection latency, or comparative performance against commercial IAM systems. Empirical validation is therefore necessary before claims concerning operational effectiveness can be generalized.

Future studies should implement the framework in controlled cloud environments and evaluate detection accuracy, policy-compliance improvement, response latency, workload disruption, and administrative cost. Such validation would transform the conceptual architecture into an empirically tested security model.

CONCLUSION

Machine identities have become foundational components of cloud-native infrastructure, yet their continuous operation, extensive privileges, and autonomous behavior create security and governance challenges that differ substantially from conventional human IAM. This research proposed an automated framework integrating identity discovery, feature-based risk assessment, least-privilege governance, lifecycle management, continuous compliance, behavioral analytics, and adaptive threat prevention.

The central contribution of the framework is the integration of these capabilities into a continuous decision loop. Instead of treating machine identities as static credentials, the model treats them as dynamic security entities whose risk changes according to privileges, behavior, lifecycle state, ownership, and compliance conditions. Computational principles derived from the supplied literature provide methodological support for feature extraction, pattern recognition, attention-based analysis, and automated classification, while Ganapathy (2025) establishes the direct relevance of automated governance for non-human identities in cloud IAM.

The framework further demonstrates that effective machine identity protection requires coordination between preventive governance and adaptive detection. Least privilege limits the potential impact of compromise, lifecycle automation reduces identity accumulation, compliance monitoring identifies governance failures, and behavioral analytics provides a mechanism for recognizing potentially compromised or misused identities.

Nevertheless, conceptual validity should not be confused with empirical validation. Future research should implement the proposed architecture in multi-cloud and cloud-native environments, evaluate its detection and response performance, and investigate explainable machine-learning approaches for identity risk scoring. Research should also examine how automated remediation can be calibrated to minimize both security exposure and disruption to legitimate production workloads.

Ultimately, machine identity security should evolve from periodic credential administration toward continuous, context-aware, and automated governance. Such an approach provides a stronger foundation for securing increasingly autonomous cloud infrastructure while maintaining accountability, compliance, and operational resilience.

REFERENCES

1. C. Agarwal and C. Bhatnagar, “Unmasking the Potential: Evaluating Image Inpainting Techniques for Masked Face Reconstruction,” Multimedia Tools and Applications, vol. 83, no. 1, pp. 893–918, 2024.

2. B. Anđić et al., “Phenomenography Study of STEM Teachers’ Conceptions of Using Three-dimensional Modeling and Printing (3DMP) in Teaching,” Journal of Science Education and Technology, vol. 32, no. 1, pp. 45–60, 2023.

3. S. Bahroun et al., “Deep 3D-LBP: CNN-based Fusion of Shape Modeling and Texture Descriptors for Accurate Face Recognition,” The Visual Computer, vol. 39, no. 1, pp. 239–254, 2023.

4. Z. Chen et al., “Transformer-based 3D Face Reconstruction with End-to-end Shape-preserved Domain Transfer,” IEEE Transactions on Circuits and Systems for Video Technology, vol. 32, no. 12, pp. 8383–8393, 2022.

5. C. Wang et al., “Cross-domain and Disentangled Face Manipulation with 3D Guidance,” IEEE Transactions on Visualization and Computer Graphics, vol. 29, no. 4, pp. 2053–2066, 2022.

6. G. D'Ettorre et al., "A Comparison Between Stereophotogrammetry and Smartphone Structured Light Technology for Three-dimensional Face Scanning," The Angle Orthodontist, vol. 92, no. 3, pp. 358–363, 2022.

7. M. C. Florkow et al., "Magnetic Resonance Imaging Versus Computed Tomography for Three-Dimensional Bone Imaging of Musculoskeletal Pathologies: A Review," Journal of Magnetic Resonance Imaging, vol. 56, no. 1, pp. 11–34, 2022.

8. S. Y. Lee et al., "In Vitro Three-dimensional (3D) Cell Culture Tools for Spheroid and Organoid Models," SLAS Discovery, vol. 28, no. 4, pp. 119–137, 2023.

9. Z. Luo et al., "A Three-dimensional Model of Student Interest During Learning Using Multimodal Fusion with Natural Sensing Technology," Interactive Learning Environments, vol. 30, no. 6, pp. 1117–1130, 2022.

10. S. C. Medin et al., "MOST-GAN: 3D Morphable StyleGAN for Disentangled Face Image Manipulation," in Proceedings of the AAAI Conference on Artificial Intelligence, vol. 36, no. 2, pp. 1962–1971, 2022.

11. M. K. Meena et al., "Partially Occluded Face Reconstruction Using Graph-based Algorithm," Journal of Electrical Engineering & Technology, vol. 19, no. 6, pp. 3655–3664, 2024.

12. N. K. Mehta et al., "Three-dimensional DenseNet Self-attention Neural Network for Automatic Detection of Student's Engagement," Applied Intelligence, vol. 52, no. 12, pp. 13803–13823, 2022.

13. H. M. U. Munir and W. S. Qureshi, "Single Image 3D Beard Face Reconstruction Approaches," International Journal of Cyber-Physical Systems (IJCPS), vol. 4, no. 1, pp. 1–17, 2022.

14. S. Pal et al., "Adapting a Swin Transformer for License Plate Number and Text Detection in Drone Images," Artificial Intelligence and Applications, vol. 1, no. 3, pp. 145–154, 2023.

15. S. Rani et al., "Three Dimensional Objects Recognition & Pattern Recognition Technique; Related Challenges: A Review," Multimedia Tools and Applications, vol. 81, no. 12, pp. 17303–17346, 2022.

16. S. Sharma and V. Kumar, "3D Face Reconstruction in Deep Learning Era: A Survey," Archives of Computational Methods in Engineering, vol. 29, no. 5, pp. 3475–3507, 2022.

17. C. Shen et al., "Large-view 3D Color Face Reconstruction from Single Image Via UV Location Map and CGAN," Journal of Computer-Aided Design & Computer Graphics, vol. 34, no. 4, pp. 614–622, 2022.

18. N. Sun et al., "3D Facial Feature Reconstruction and Learning Network for Facial Expression Recognition in the Wild," IEEE Transactions on Cognitive and Developmental Systems, vol. 15, no. 1, pp. 298–309, 2022.

19. E. Tretschk et al., "State of the Art in Dense Monocular Non-rigid 3D Reconstruction," Computer Graphics Forum, vol. 42, no. 2, pp. 485–520, 2023.

20. Z. Zhou et al., "Replay Attention and Data Augmentation Network for 3D Face and Object Reconstruction," IEEE Transactions on Biometrics, Behavior and Identity Science, vol. 5, no. 3, pp. 308–320, 2023.

21. 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