The American Journal of Engineering and Technology
ISSN 2689-0984 Volume 08 - 2026

Continuous Identity Governance and Automated Threat Protection for Non-Human Identities in Multi-Cloud IAM

Kwame Mensah
Department of Computer Science, Ghana Institute of Digital Technology, Ghana

Abena Owusu
Department of Computer Science, Ghana Institute of Digital Technology, Ghana

Received: 05 Jun 2026 | Received Revised Version: 15 Jui 2026 | Accepted: 25 Aug 2026 | Published: 03 Sep 2026

Volume 08 Issue 09 2026 | Crossref DOI: 10.37547/tajet/Volume08Issue09-03

Abstract

The rapid adoption of cloud-native applications, distributed services, automated pipelines, containers, workload orchestration, and machine-to-machine communication has resulted in a substantial expansion of non-human identities (NHIs) within enterprise environments. Unlike conventional human identities, NHIs frequently operate continuously, interact programmatically, and require access across multiple cloud platforms, creating governance and security challenges that cannot be adequately addressed through static access-control practices. This research develops a conceptual framework for continuous identity governance and automated threat protection for NHIs in multi-cloud Identity and Access Management (IAM). The proposed approach combines identity discovery, contextual authorization, behavioral monitoring, graph-oriented relationship analysis, risk evaluation, adaptive policy enforcement, and automated remediation. The methodology synthesizes the computational principles evident in the provided literature, particularly distributed graph processing, parallel computation, clustering, data mining, and scalable analytics, and applies them conceptually to NHI governance. The analysis indicates that graph-based identity representations can improve visibility into relationships among workloads, credentials, resources, and permissions, while distributed processing can support continuous analysis at multi-cloud scale. The study further argues that automated governance should be risk-adaptive rather than exclusively rule-based, allowing access privileges and protective controls to respond dynamically to changing identity behavior. The proposed framework contributes a research-oriented foundation for continuous NHI governance while recognizing limitations related to heterogeneous cloud environments, behavioral baselines, false positives, policy conflicts, and computational overhead.

Keywords: Non-Human Identities, Multi-Cloud IAM, Continuous Identity Governance, Automated Threat Protection, Identity Graphs, Adaptive Access Control, Cloud Security, Behavioral Analytics, Distributed Computing

Cite This Article: Mensah, K., & Owusu, A. (2026). Continuous Identity Governance and Automated Threat Protection for Non-Human Identities in Multi-Cloud IAM. The American Journal of Engineering and Technology, 8(09), 15–22. Retrieved from https://theamericanjournals.com/index.php/tajet/article/view/8362

1. Introduction

Cloud computing has transformed enterprise information systems from relatively centralized environments into distributed ecosystems containing applications, containers, microservices, APIs, automated workloads, data pipelines, and machine-to-machine communication. Within these environments, access is no longer exclusively performed by employees. Service accounts, application identities, workload identities, API credentials, automation agents, containers, and other machine-operated principals increasingly perform operational functions that previously depended on human interaction. These entities constitute non-human identities (NHIs), and their growing number introduces a distinct identity-governance problem.

The central security difficulty is not simply the existence of NHIs but their persistent and programmatic access to resources. An NHI may authenticate without human

intervention, invoke multiple services, access distributed datasets, and operate across several cloud environments. Consequently, an identity that was legitimate when created may subsequently become excessive, unused, misconfigured, compromised, or inconsistent with its intended workload. Traditional periodic access reviews therefore provide limited protection against continuously changing NHI behavior.

The problem becomes more complex in multi-cloud environments. Different cloud platforms may employ different identity models, policy languages, logging mechanisms, resource structures, and privilege semantics. A governance mechanism that evaluates identities independently within each cloud may fail to recognize cross-cloud relationships. For example, an automation identity could authenticate to a workload in one cloud, retrieve a credential through another service, and access a storage resource hosted elsewhere. Each individual event might appear legitimate, while the combined sequence could represent excessive privilege or anomalous activity.

The scalability requirements of this problem are closely related to challenges addressed in distributed computing and large-scale graph analytics. Research on parallel graph algorithms demonstrates that graph computations can be designed to remain efficient and scalable even when operating over large datasets (Dhulipala et al., 2021). Similarly, large-scale graph engines and transactional graph data structures demonstrate the importance of efficient processing and representation when relationships become numerous and dynamic (Fan et al., 2021; Fuchs et al., 2022).

The objective of this research is therefore to develop a conceptual framework in which NHI governance becomes continuous, contextual, graph-aware, and automated. Rather than treating identity governance as a periodic administrative activity, the proposed approach considers identity security as an ongoing computational process. This perspective is consistent with the broader findings of scalable data-processing research, where parallelization, partitioning, clustering, and distributed execution are necessary to process increasingly complex datasets efficiently (Dafir et al., 2021; Kaur et al., 2022).

The study has four principal objectives. First, it conceptualizes a unified representation of NHIs and their relationships across multiple cloud environments. Second, it proposes a continuous risk-analysis mechanism based on identity attributes, permissions, resource relationships, and observed behavior. Third, it develops an automated protection model capable of adapting authorization and initiating remediation actions. Fourth, it examines the theoretical and practical limitations of applying distributed analytics and graph-oriented techniques to multi-cloud identity governance.

The significance of this work lies in connecting identity security with scalable computational techniques. Ganapathy (2025) specifically emphasizes automated governance and protection of NHIs in cloud IAM, providing a direct conceptual foundation for treating NHI governance as an automated security function rather than merely a credential-management activity. The present research extends that perspective toward a multi-cloud and continuously adaptive governance architecture.

2. Literature Review

2.1 Distributed Analytics as a Foundation for Identity Governance

Large-scale identity governance requires processing heterogeneous and continuously changing information. Identity inventories, permissions, authentication events, workload relationships, resource metadata, and policy decisions collectively form a high-dimensional analytical dataset. Research on big-data processing indicates that parallel and distributed approaches are increasingly necessary when computational workloads exceed the practical capabilities of centralized processing (Kumar and Mohbey, 2022).

Dhulipala et al. (2021) demonstrate that theoretically efficient parallel graph algorithms can achieve practical scalability. This is particularly relevant to NHI governance because identity relationships naturally form graphs: identities connect to applications, applications connect to resources, credentials connect to workloads, and permissions connect principals to actions. Efficient graph processing can consequently support continuous evaluation without requiring the entire identity environment to be analyzed sequentially.

Fan et al. (2021) present GraphScope as a unified engine for big graph processing, reinforcing the potential value of unified graph-processing architectures. For multi-cloud IAM, such a principle suggests that identity relationships should be analyzed through a common logical model even when their underlying cloud implementations differ.

2.2 Graph-Based Representation of Identity

Relationships

The identity environment can be modeled as a heterogeneous graph G=(V,E)G=(V,E), where vertices represent NHIs, workloads, services, resources, credentials, policies, and cloud accounts, while edges represent relationships such as authenticates-to, invokes, reads, writes, assumes, and delegates.

This representation provides analytical advantages over isolated identity records. A credential may appear legitimate when examined independently, but its graph neighborhood may reveal connections to resources that fall outside its intended operational function. Bouhenni et al. (2021), through their examination of distributed graph pattern matching, demonstrate the importance of identifying structural patterns in massive graphs. In NHI security, analogous pattern matching could identify suspicious permission chains, unexpected trust relationships, or unusual cross-service access paths.

Czumaj et al. (2021) further demonstrate the relevance of graph sparsification for massively parallel computation. Applied conceptually to IAM, graph reduction could remove redundant or low-value relationships while retaining security-relevant paths. This may reduce computational requirements for continuous analysis, although excessive simplification could eliminate relationships necessary for accurate risk evaluation.

Fuchs et al. (2022) address transactional graph data structures, highlighting another important requirement: identity relationships are not static. Policies, workloads, credentials, and resource relationships change continuously. A governance architecture must therefore support dynamic updates rather than relying solely on periodically generated identity graphs.

2.3 Clustering and Behavioral Analysis

NHI security requires more than structural analysis. Legitimate workloads often produce predictable behavioral patterns, whereas compromised identities may generate deviations. Clustering techniques can support the grouping of identities according to behavioral or structural characteristics.

Dafir et al. (2021) survey parallel clustering algorithms for big data, demonstrating the relevance of scalable clustering for large datasets. In a multi-cloud NHI environment, clustering could group identities according to access frequency, resource types, invocation patterns, geographical or temporal characteristics, and privilege structures.

For example, a deployment service that normally accesses container registries and compute resources during predefined deployment windows could establish a behavioral cluster. A sudden transition toward unrelated databases or administrative APIs would represent a measurable deviation from the cluster profile. Such analysis should not automatically imply compromise, because legitimate operational changes also generate behavioral deviations. Therefore, behavioral analytics should contribute to a risk score rather than act as an isolated decision mechanism.

2.4 Parallel Scheduling and Continuous Governance

Continuous governance creates a scheduling problem. Identity analysis, policy evaluation, anomaly detection, and remediation must compete for computational resources. Kaur et al. (2022) investigate multi-level parallel scheduling of dependent tasks using graph partitioning and hybrid approaches over edge-cloud environments. Their work suggests that identity analytics can similarly be divided into dependent and independent processing stages.

For instance, identity discovery can occur concurrently with behavioral feature extraction, while graph-based risk analysis can subsequently integrate their outputs. Such decomposition can reduce latency while preserving analytical dependencies.

Hoang et al. (2021) examine a streaming edge partitioner for distributed graph analytics. This is particularly relevant because identity telemetry is inherently streaming. Rather than processing all events as a single batch, security systems can partition incoming authentication and authorization events across analytical workers.

2.5 Data Mining and Intelligent Security Analytics

Gheisari et al. (2023) review data-mining techniques for web mining, illustrating how large datasets can be transformed into useful patterns and knowledge. The same general analytical principle can be applied to NHI telemetry: raw authentication and authorization events must be transformed into behavioral indicators and risk-relevant patterns.

Jain et al. (2021) demonstrate the application of deep learning to qualitative user-generated content, while Muthu et al. (2021) examine deep-learning-based text summarization. Although these studies address different

domains, they illustrate a broader methodological principle: complex raw data can be transformed through machine-learning pipelines into higher-level representations. For NHI governance, this principle supports the transformation of raw logs into behavioral features rather than treating individual events as isolated security signals.

Ibtisum et al. (2023) compare MapReduce and Apache Spark paradigms for big-data processing, highlighting the importance of choosing appropriate processing architectures according to workload requirements. In continuous identity governance, stream-oriented and distributed processing architectures may be preferable when decisions must be generated rapidly.

2.6 Supporting Computational Evidence

Several provided studies address specialized applications of machine learning and distributed processing. Kasinathan et al. (2021) demonstrate machine-learning techniques for automated classification and detection, while Ngugi et al. (2021) examine automated recognition through image-processing methods. Wu et al. (2021) discuss parallel and distributed computing for remotely sensed big data, and Yin and Shi (2022) compare big-data computing with high-performance computing.

Although these studies do not directly investigate IAM, collectively they reinforce the technical premise that large-scale automated analysis depends on appropriate computational architectures. Kumar and Mohbey (2022) similarly establish the broader importance of parallel and distributed approaches for pattern-mining workloads.

The literature gap is therefore identifiable: the provided computational research offers substantial foundations for scalable analytics, graph processing, pattern matching, clustering, and distributed computation, but these principles are not directly integrated into a continuous multi-cloud NHI governance architecture. Ganapathy (2025) provides the most direct NHI governance foundation, while this research positions that problem within a broader scalable computational model.

3. Methodology

3.1 Research Design

This study follows a conceptual research-and-review methodology. Rather than evaluating a commercial IAM platform, it synthesizes the provided literature into an integrated framework for continuous NHI governance. The methodology consists of five logical layers: identity discovery, relationship modeling, continuous risk analytics, adaptive policy enforcement, and automated remediation.

The architecture assumes that multi-cloud environments generate heterogeneous identity and security telemetry. These data are normalized into a common identity model before analytical processing. The conceptual design draws on distributed processing principles discussed by Kumar and Mohbey (2022), graph-processing approaches discussed by Fan et al. (2021), and scalable parallel algorithms examined by Dhulipala et al. (2021).

3.2 Identity Discovery and Normalization

The first stage establishes an inventory of NHIs. Each identity record should conceptually contain an identifier, identity type, owning application, associated workload, credential mechanism, permissions, cloud environment, resource relationships, creation information, and observed activity.

The critical methodological principle is normalization. A service identity represented differently in two cloud providers should nevertheless be mapped to a common logical identity schema. Without normalization, cross-cloud analysis becomes fragmented.

Consider an organization operating workloads in three clouds. An automated deployment identity may have different provider-specific representations, but the governance engine should recognize that these representations correspond to the same operational function. This abstraction enables cross-cloud risk analysis.

3.3 Identity Relationship Graph

After normalization, identity entities are represented as a graph. Nodes may represent identities, applications, services, resources, policies, credentials, and cloud accounts. Edges encode relationships such as authentication, authorization, delegation, invocation, ownership, and resource access.

The graph should support incremental updates because identity relationships change continuously. Fuchs et al. (2022) provide relevant conceptual support for transactional graph structures, while Hoang et al. (2021) highlight streaming-oriented graph partitioning.

A simplified security relationship could be represented as:

NHI → Workload → Service → Resource → Permission

A suspicious access path might therefore be detected not because any single relationship is invalid, but because the complete path violates the expected operational profile.

3.4 Continuous Risk Evaluation

The proposed framework assigns a dynamic risk value to each NHI. Conceptually, risk can be modeled as a function:

Ri=f(Pi,Bi,Gi,Ti,Ci)R_i = f(P_i, B_i, G_i, T_i, C_i)

where PiP_i represents privilege characteristics, BiB_i represents behavioral deviation, GiG_i represents graph relationships, TiT_i represents temporal activity, and CiC_i represents contextual conditions.

This is intentionally a conceptual model rather than a fixed production formula. The objective is to avoid binary decisions in which an identity is simply classified as trusted or untrusted.

For example, an NHI with broad administrative privileges may initially receive elevated risk. If it subsequently accesses an unusual resource outside its established graph neighborhood, the risk value should increase further. Conversely, a legitimate emergency operation approved through appropriate governance context could reduce the significance of that anomaly.

Clustering can support behavioral baselines, consistent with the scalable clustering approaches examined by Dafir et al. (2021). Graph pattern matching can identify structurally unusual access paths, drawing conceptually on Bouhenni et al. (2021).

3.5 Adaptive Policy Enforcement

Risk evaluation must lead to an actionable security decision. The framework therefore proposes several policy responses ranging from continued access to additional verification, privilege reduction, temporary suspension, credential rotation, or complete isolation.

The important distinction is that policy enforcement is adaptive. A static policy might permanently grant an identity access to a database. A continuous governance system instead evaluates whether that access remains justified under current conditions.

Ganapathy (2025) directly supports the conceptual movement toward automated governance and protection of NHIs. The proposed framework extends this principle by incorporating graph relationships and multi-cloud analytics into continuous decision-making.

3.6 Automated Remediation

The final stage translates risk decisions into automated controls. Possible actions include temporary privilege reduction, disabling unused credentials, restricting resource access, forcing credential rotation, isolating a workload, or initiating an investigation workflow.

Automation must nevertheless incorporate safeguards. An incorrect automated response could disrupt production systems. Therefore, high-impact actions should require stronger evidence than low-impact controls. A low-risk anomaly may produce additional monitoring, whereas a high-confidence compromise indicator may justify immediate containment.

Parallel execution can improve remediation responsiveness. Kaur et al. (2022) provide relevant evidence that dependent computational tasks can be scheduled through graph-aware approaches, while Ibtisum et al. (2023) illustrate the broader importance of selecting suitable distributed processing paradigms.

3.7 Scalability Architecture

The architecture should separate ingestion, analytical processing, policy evaluation, and remediation. Incoming telemetry can be distributed across processing nodes, while graph partitions can enable parallel analysis.

Czumaj et al. (2021) demonstrate the relevance of graph sparsification for massively parallel computation. In the proposed framework, graph optimization could prioritize security-relevant relationships while reducing computational overhead. However, sparsification must preserve critical privilege and trust paths.

The framework consequently favors a hybrid architecture in which frequently changing security signals are processed incrementally while deeper relationship analysis is performed periodically or when elevated risk conditions occur.

4. Results

The conceptual synthesis produces five principal findings.

First, continuous NHI governance is fundamentally a scalability problem. The number of identities, permissions, workloads, resources, and authorization

events can grow beyond the practical limits of manual review. The distributed-processing literature demonstrates that large computational workloads benefit from parallel and scalable architectures (Kumar and Mohbey, 2022; Dhulipala et al., 2021). Therefore, continuous governance requires computational infrastructure capable of processing identity telemetry without excessive decision latency.

Second, graph representation provides stronger contextual visibility than isolated identity records. An identity's security significance depends on its relationships with workloads, resources, credentials, and policies. Graph processing can expose multi-step access paths that conventional identity inventories may overlook. Distributed graph pattern matching and graph analytics provide a suitable theoretical foundation for this capability (Bouhenni et al., 2021; Fan et al., 2021).

Third, behavioral analytics can complement deterministic IAM policies. Static authorization determines whether access is technically permitted, but behavioral analysis evaluates whether that access is consistent with expected operational activity. Clustering and pattern-mining approaches provide a basis for establishing behavioral groups and identifying deviations (Dafir et al., 2021; Gheisari et al., 2023).

Fourth, continuous governance should be risk-adaptive rather than uniformly restrictive. The framework indicates that binary trust models are insufficient for dynamic workloads. An NHI should be evaluated according to privilege, behavior, context, and graph relationships. This supports differentiated responses, from monitoring to privilege restriction and automated containment. The emphasis on automated NHI governance is also consistent with Ganapathy (2025).

Fifth, automation creates a governance trade-off between response speed and operational safety. Immediate automated remediation can reduce exposure during a genuine compromise but can also interrupt legitimate workloads when anomalies are incorrectly interpreted. Consequently, automated responses should be proportional to confidence and potential impact. The computational literature supports scalable processing, but scalability alone does not guarantee accurate security decisions (Kaur et al., 2022; Ibtisum et al., 2023).

Overall, the findings indicate that the most promising architecture is a layered system combining normalized identity data, dynamic identity graphs, parallel analytics, behavioral modeling, contextual risk scoring, and graduated remediation. This architecture is more suitable for multi-cloud environments than isolated periodic identity reviews because it treats identity state as continuously changing.

5. Discussion

The findings suggest a conceptual shift from identity administration to identity intelligence. Traditional governance primarily answers whether an identity possesses authorization. Continuous governance additionally asks whether the authorization remains appropriate, whether the identity is behaving consistently with its purpose, and whether its current relationships introduce unexpected risk.

The theoretical contribution of the framework is the integration of identity governance with scalable graph and distributed-computing principles. The graph abstraction allows IAM relationships to be treated as computational structures rather than static configuration records. This is particularly valuable because the security meaning of an NHI often emerges from combinations of relationships rather than from individual permissions. Bouhenni et al. (2021) and Fan et al. (2021) provide computational foundations for analyzing large relationship structures, while Fuchs et al. (2022) reinforce the importance of dynamic graph representations.

A second implication concerns scalability. Continuous monitoring can generate substantially more analytical workload than periodic auditing. Parallel graph algorithms, partitioning, clustering, and distributed processing therefore become architectural requirements rather than optional optimizations. Dhulipala et al. (2021), Czumaj et al. (2021), and Kaur et al. (2022) collectively indicate why workload decomposition and efficient parallel execution are important when analytical datasets become large.

However, automation introduces significant limitations. The first is false-positive risk. Legitimate infrastructure changes can resemble malicious behavior. A deployment identity may suddenly access a new resource because an application was legitimately redesigned. If behavioral deviation automatically results in account suspension, governance can become an operational liability.

The second limitation is cross-cloud semantic inconsistency. Normalizing identities across cloud providers does not eliminate differences in policy

semantics, resource structures, or privilege models. A unified analytical layer may therefore require abstractions that inevitably lose some provider-specific information.

The third limitation is graph complexity. Complete identity graphs may become extremely dense. Graph sparsification can improve efficiency, as indicated by Czumaj et al. (2021), but removing apparently insignificant edges can potentially hide attack paths. Security-oriented graph reduction therefore requires careful preservation criteria.

The fourth limitation concerns behavioral baselines. A newly deployed workload has limited historical behavior, making anomaly detection less reliable. Similarly, seasonal or emergency workloads may legitimately violate established patterns. Clustering and machine-learning techniques can assist, but their outputs should remain contextual risk indicators rather than unquestionable security decisions.

The fifth limitation is computational cost. Continuous processing, graph updates, clustering, and risk scoring require substantial infrastructure. Ibtisum et al. (2023) demonstrate that processing paradigms differ in their suitability for particular workloads, suggesting that a single processing model may not optimally support every component of continuous identity governance.

Practically, the framework therefore recommends graduated automation. Low-confidence anomalies should trigger monitoring and additional analysis. Moderate-confidence risks may result in privilege reduction or additional authorization requirements. High-confidence risks can justify automated containment. This approach preserves the principal advantage of automation—rapid response—while reducing the operational risks associated with indiscriminate remediation.

The discussion also reinforces the importance of treating NHI governance as a continuous lifecycle. Identity creation, privilege assignment, behavioral monitoring, risk reassessment, remediation, and retirement should not be independent administrative activities. They should form an integrated feedback loop. Ganapathy (2025) provides a direct foundation for this automated governance perspective, while the distributed analytics literature supplies the computational mechanisms necessary to make continuous analysis feasible at scale.

6. Conclusion

This research presented a conceptual framework for continuous identity governance and automated threat protection for non-human identities in multi-cloud IAM. The study identified that NHI security cannot be effectively addressed through static permissions or periodic access reviews alone because machine identities operate continuously, interact with distributed resources, and may span multiple cloud environments.

The proposed framework combines five capabilities: identity discovery and normalization, graph-based relationship modeling, continuous risk evaluation, adaptive policy enforcement, and automated remediation. The literature synthesis demonstrates that distributed graph processing, parallel algorithms, clustering, streaming analytics, and scalable data-processing techniques provide relevant theoretical foundations for implementing these capabilities.

The principal contribution is the positioning of NHI governance as a continuous computational security process. Identity state is modeled as dynamic, behavioral evidence is incorporated into risk assessment, and authorization decisions can adapt according to contextual changes. This perspective can provide stronger visibility into excessive privilege, unusual access paths, and anomalous workload behavior than isolated identity inventories.

Nevertheless, the framework remains conceptual and should be validated experimentally. Future research should develop a prototype across multiple cloud platforms, establish standardized NHI datasets, evaluate graph-based risk models, quantify false-positive and false-negative rates, and measure the computational cost of continuous analysis. Further work should also investigate explainable risk scoring, policy-conflict resolution, privacy-preserving telemetry processing, and safe automated remediation.

Ultimately, effective multi-cloud NHI protection requires a balance between continuous visibility, analytical scalability, adaptive authorization, and operational safety. The integration of identity governance with distributed computational intelligence provides a promising foundation for achieving that balance.

References

1. M. Babar et al., “An optimized IoT-enabled big data analytics architecture for edge-cloud computing,” *IEEE Internet of Things Journal*, vol. 10, no. 5, pp. 3995–4005, 2022.

2. S. Bouhenni et al., "A survey on distributed graph pattern matching in massive graphs," *ACM Computing Surveys (CSUR)*, vol. 54, no. 2, pp. 1–35, 2021.
3. A. Czumaj et al., "Graph sparsification for derandomizing massively parallel computation with low space," *ACM Transactions on Algorithms (TALG)*, vol. 17, no. 2, pp. 1–27, 2021.
4. Z. Dafir et al., "A survey on parallel clustering algorithms for big data," *Artificial Intelligence Review*, vol. 54, no. 4, pp. 2411–2443, 2021.
5. L. Dhulipala et al., "Theoretically efficient parallel graph algorithms can be fast and scalable," *ACM Transactions on Parallel Computing (TOPC)*, vol. 8, no. 1, pp. 1–70, 2021.
6. W. Fan et al., "GraphScope: A unified engine for big graph processing," *Proceedings of the VLDB Endowment*, vol. 14, no. 12, pp. 2879–2892.
7. P. Fuchs et al., "Sortledton: A universal, transactional graph data structure," *Proceedings of the VLDB Endowment*, vol. 15, no. 6, pp. 1173–1186.
8. M. Gheisari et al., "Data mining techniques for web mining: A survey," *Artificial Intelligence and Applications*, vol. 1, no. 1, pp. 3–10, 2023.
9. L. Hoang et al., "Cusp: A customizable streaming edge partitioner for distributed graph analytics," *ACM SIGOPS Operating Systems Review*, vol. 55, no. 1, pp. 47–60, 2021.
10. G. Hou et al., "Massively parallel algorithms for personalized PageRank," *Proceedings of the VLDB Endowment*, vol. 14, no. 9, pp. 1668–1680.
11. S. Ibtisum et al., "A comparative analysis of big data processing paradigms: MapReduce vs. Apache Spark," *World Journal of Advanced Research and Reviews*, vol. 20, no. 1, pp. 1089–1098.
12. P. K. Jain et al., "A hybrid CNN-LSTM: A deep learning approach for consumer sentiment analysis using qualitative user-generated contents," *Transactions on Asian and Low-Resource Language Information Processing*, vol. 20, no. 5, pp. 1–15, 2021.
13. T. Kasinathan et al., "Insect classification and detection in field crops using modern machine learning techniques," *Information Processing in Agriculture*, vol. 8, no. 3, pp. 446–457, 2021.
14. M. Kaur et al., "Multi-level parallel scheduling of dependent-tasks using graph-partitioning and hybrid approaches over edge-cloud," *Soft Computing*, vol. 26, no. 11, pp. 5347–5362, 2022.
15. A. Kochsiek and R. Gemulla, "Parallel training of knowledge graph embedding models: A comparison of techniques," *Proceedings of the VLDB Endowment*, vol. 15, no. 3, pp. 633–645, 2021.
16. S. Kumar and K. K. Mohbey, "A review on big data based parallel and distributed approaches of pattern mining," *Journal of King Saud University-Computer and Information Sciences*, vol. 34, no. 5, pp. 1639–1662, 2022.
17. B. Muthu et al., "A framework for extractive text summarization based on deep learning modified neural network classifier," *Transactions on Asian and Low-Resource Language Information Processing*, vol. 20, no. 3, pp. 1–20, 2021.
18. L. C. Ngugi et al., "Recent advances in image processing techniques for automated leaf pest and disease recognition—A review," *Information Processing in Agriculture*, vol. 8, no. 1, pp. 27–51, 2021.
19. Z. Wu et al., "Recent developments in parallel and distributed computing for remotely sensed big data processing," *Proceedings of the IEEE*, vol. 109, no. 8, pp. 1282–1305, 2021.
20. F. Yin and F. Shi, "A comparative survey of big data computing and HPC: From a parallel programming model to a cluster architecture," *International Journal of Parallel Programming*, vol. 50, no. 1, pp. 27–64, 2022.
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.