International Journal of Emerging Trends in Engineering and Management Research (IJETEMR)

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104007

Autonomous Multi-Cloud Governance through Intelligent Agentic AI and Dynamic Policy Enforcement

Pavan Srikanth Subba Raju Patchamatla

Cloud Application Engineer, RK Infotech LLC, USA

Publication History: Received: 05.07.2026; Revised: 02.08.2026; Accepted: 07.08.2026; Published: 20.08.2026

ABSTRACT: Multi-cloud environments have become increasingly important for organizations seeking scalability, resilience, cost optimization, and flexibility across heterogeneous cloud providers. However, managing governance across multiple cloud platforms introduces significant challenges related to security, compliance, configuration consistency, identity management, resource optimization, and policy enforcement. Traditional governance approaches largely depend on static rules, manually configured controls, and centralized monitoring mechanisms, which may be inadequate for highly dynamic cloud infrastructures. This paper proposes an autonomous multi-cloud governance framework based on intelligent agentic artificial intelligence and dynamic policy enforcement. The proposed approach employs autonomous AI agents capable of continuously observing cloud environments, interpreting organizational policies, identifying deviations, reasoning about appropriate responses, and executing corrective actions across heterogeneous cloud platforms. A dynamic policy layer translates governance requirements into machine-enforceable controls while adapting enforcement decisions according to contextual information such as workload characteristics, risk levels, regulatory requirements, and operational conditions. The research methodology combines architectural design, scenario-based experimentation, quantitative evaluation, and comparative analysis. Key evaluation dimensions include policy-compliance accuracy, detection latency, remediation effectiveness, false-positive rates, resource overhead, and cross-cloud consistency. The proposed framework aims to establish a closed governance loop in which perception, reasoning, policy evaluation, action, and continuous learning operate together. Such an approach can strengthen governance automation while reducing administrative effort and improving the responsiveness of multi-cloud security and compliance management.

KEYWORDS: Multi-cloud governance, agentic AI, autonomous agents, dynamic policy enforcement, cloud security, policy automation, artificial intelligence, compliance management, cloud orchestration, intelligent automation

I. INTRODUCTION

The adoption of multiple cloud service providers has transformed enterprise information technology by allowing organizations to distribute workloads according to requirements involving performance, availability, geographic coverage, cost, specialized services, and regulatory constraints. Instead of relying exclusively on a single cloud provider, organizations increasingly operate heterogeneous environments incorporating infrastructure, platforms, applications, containers, databases, and serverless services from different providers. Although multi-cloud architectures provide operational flexibility, they also create substantial governance complexity. Each provider exposes different management interfaces, identity mechanisms, security configurations, resource models, monitoring facilities, and policy controls. Consequently, maintaining consistent organizational governance across these environments can become difficult when policies are implemented through provider-specific tools and manually coordinated operational procedures. Conventional governance mechanisms generally depend on predefined rules, periodic audits, centralized dashboards, and human intervention. These approaches can identify policy violations, but they may not respond rapidly enough to continuously changing cloud workloads and configurations. The increasing scale and dynamism of cloud infrastructures therefore create a requirement for governance mechanisms capable of continuously observing environments, reasoning about policy requirements, and responding autonomously to deviations.

Agentic artificial intelligence provides a potential foundation for addressing this challenge because intelligent agents can combine perception, reasoning, planning, tool utilization, and action within an ongoing operational loop. In a multi-cloud governance context, agents can monitor resources and configurations, interpret organizational policies, evaluate

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104007

contextual risks, and initiate appropriate remediation activities. Rather than treating governance as a static compliance-checking process, an agentic framework can conceptualize governance as a continuous cycle of observation, analysis, decision-making, enforcement, and verification. Dynamic policy enforcement further strengthens this approach by allowing governance decisions to consider contextual factors rather than applying identical controls to every workload. For example, an enforcement mechanism could distinguish between production and development resources, evaluate data sensitivity, consider regulatory requirements associated with a geographic region, and modify enforcement actions according to the assessed level of risk. The combination of intelligent agents and dynamic policies consequently provides an opportunity to develop autonomous governance capable of coordinating heterogeneous cloud environments while preserving organizational control. This research investigates such an approach by developing a conceptual architecture and methodology for autonomous multi-cloud governance, with particular emphasis on policy consistency, adaptive enforcement, security, compliance, remediation effectiveness, and operational efficiency.

II. LITERATURE REVIEW

Cloud governance research has traditionally focused on establishing structured mechanisms for controlling cloud resources, identities, security configurations, costs, compliance requirements, and operational processes. Policy-based management has emerged as an important approach because organizational requirements can be expressed as formal rules that determine whether particular cloud configurations or activities are permitted. Infrastructure-as-code and policy-as-code have further supported governance automation by allowing infrastructure configurations and compliance requirements to be represented in machine-readable formats. However, multi-cloud environments complicate policy management because cloud providers expose different resource abstractions, APIs, identity models, and security capabilities. A governance rule expressed directly against one provider's resources may not transfer automatically to another provider. Existing cloud-management approaches therefore frequently require abstraction layers, centralized management platforms, or provider-specific policy implementations. Although these approaches can improve consistency, they may remain dependent on predefined conditions and deterministic workflows. Moreover, many governance mechanisms emphasize detection and reporting rather than autonomous reasoning and remediation. This limitation becomes particularly relevant in environments where resources are continuously created, modified, migrated, scaled, or terminated. Effective governance consequently requires mechanisms that can interpret policy requirements within changing operational contexts and coordinate enforcement across heterogeneous infrastructures.

Artificial intelligence and autonomous-agent research offers several concepts that can extend conventional cloud governance. Intelligent agents can perceive environmental information, maintain contextual state, reason about objectives, and perform actions through available tools or interfaces. In cloud computing, AI has been investigated for resource management, anomaly detection, workload optimization, security monitoring, and operational automation. More recent agentic AI approaches extend these capabilities by enabling systems to decompose objectives, select tools, execute multi-step tasks, and evaluate outcomes. For governance, these capabilities create the possibility of agents that move beyond simple anomaly detection toward autonomous policy interpretation and remediation. Nevertheless, agentic autonomy introduces important governance concerns of its own. An agent operating with extensive cloud permissions could potentially execute inappropriate changes if its reasoning is incorrect, its contextual information is incomplete, or its policy interpretation is ambiguous. Research on trustworthy AI therefore highlights the importance of explainability, authorization boundaries, auditability, human oversight, and controlled autonomy. A governance architecture must consequently ensure that intelligent agents themselves operate under explicit constraints. Dynamic policy enforcement can provide such constraints by defining allowable actions, approval requirements, escalation conditions, and risk thresholds. The literature thus indicates a need to combine the adaptability of intelligent agents with deterministic governance controls rather than treating autonomous AI as an unrestricted decision-making mechanism.

Another important research area concerns continuous compliance and adaptive security in distributed cloud environments. Traditional compliance assessments are often periodic, meaning that a configuration can become non-compliant between assessment cycles. Continuous monitoring addresses this limitation by evaluating cloud resources and activities as they change. However, monitoring alone does not guarantee effective remediation. A mature governance architecture requires an integrated mechanism that can detect a violation, determine its significance, identify an appropriate response, implement the response, and verify whether the resulting state satisfies the relevant policy. This suggests a closed-loop model in which policy evaluation and operational enforcement are continuously connected. Existing research also indicates that automated remediation must account for potential service disruption, dependencies, and competing policies. For example, immediately modifying a security configuration could restore

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104007

compliance while simultaneously affecting application availability. An intelligent governance framework should therefore evaluate policy priorities, resource dependencies, risk levels, and operational context before executing corrective actions. The literature collectively supports the development of an autonomous multi-cloud governance model that integrates policy-as-code, continuous monitoring, intelligent agents, adaptive risk evaluation, and controlled remediation while retaining auditable decision processes and appropriate human intervention.

III. RESEARCH METHODOLOGY

This research adopts a design-oriented methodology to develop and evaluate an autonomous multi-cloud governance framework based on intelligent agentic AI and dynamic policy enforcement. The methodology begins with the identification of governance requirements associated with heterogeneous cloud environments, including security, compliance, identity and access management, resource configuration, data protection, operational consistency, and auditability. These requirements are transformed into formal governance policies that can be represented through a policy-as-code mechanism. The proposed architecture consists of several logically connected components: a multi-cloud observation layer, a normalized cloud-resource representation layer, a policy repository, an agentic reasoning layer, a dynamic policy evaluation engine, an enforcement and orchestration layer, and an audit and feedback layer. The observation layer collects information from participating cloud platforms through provider APIs, monitoring interfaces, configuration-management systems, and security telemetry. Because cloud providers expose different resource models, the collected information is transformed into a common representation so that governance agents can reason about resources independently of provider-specific implementation details. The policy repository stores organizational requirements in machine-readable form and associates each policy with conditions, priorities, permitted actions, exceptions, and escalation requirements. The agentic reasoning layer is responsible for interpreting observations, identifying potentially non-compliant states, gathering additional contextual information, determining candidate responses, and invoking authorized tools. The dynamic policy engine evaluates proposed actions against current organizational constraints before enforcement occurs. This separation between reasoning and authorization is important because an intelligent agent may propose an action, while the policy engine determines whether that action is permissible. The enforcement layer then communicates with cloud-provider APIs or orchestration mechanisms to implement approved changes. Finally, the audit layer records observations, policy decisions, agent reasoning metadata, actions, outcomes, and subsequent verification results, creating a traceable governance history. The overall methodology therefore models governance as a continuous feedback loop rather than a periodic assessment activity.

The experimental environment will employ a representative multi-cloud testbed containing heterogeneous workloads distributed across at least two logically distinct cloud environments or cloud-simulation environments. The testbed will include virtual machines, containerized workloads, storage resources, identities, network configurations, and application services so that the framework can be evaluated against different categories of governance requirements. A collection of policy scenarios will be developed to represent common governance conditions, including unrestricted public storage, excessive identity privileges, missing encryption configurations, non-compliant geographic placement, insecure network exposure, insufficient logging, and unauthorized resource changes. Each scenario will contain a defined initial state, applicable policy, expected compliance condition, and acceptable remediation outcome. The agentic system will continuously receive resource-state information and event notifications and will determine whether the observed state satisfies the relevant policies. When a deviation is identified, the agent will construct a response plan using contextual information such as resource criticality, environment type, policy priority, dependency relationships, and risk classification. Candidate actions may include modifying a configuration, restricting access, enabling a required security control, isolating a resource, creating an alert, or escalating the event for human approval. To prevent uncontrolled autonomy, enforcement actions will be divided into predefined authorization categories. Low-risk actions may be automatically executed, medium-risk actions may require additional policy validation, and high-impact actions may require explicit human approval. This experimental design allows the research to evaluate autonomy without assuming that every cloud-management decision should be delegated entirely to an AI agent. Each remediation cycle will also include post-action verification to determine whether the target resource has returned to a compliant state.

The evaluation will use quantitative metrics designed to measure both governance effectiveness and operational cost. Policy-compliance accuracy will measure the proportion of correctly identified compliant and non-compliant states, while remediation success rate will measure the proportion of detected violations for which the system produces an acceptable corrective outcome. Detection latency will measure the time between the occurrence of a policy violation and its identification by the governance system, whereas remediation latency will measure the time between detection and successful enforcement. False-positive and false-negative rates will be calculated to determine the reliability of

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104007

policy evaluation and agentic decision-making. Cross-cloud policy consistency will assess whether equivalent governance requirements produce equivalent outcomes across heterogeneous cloud providers despite differences in underlying implementation. Resource overhead will measure the computational and network resources consumed by continuous monitoring, policy evaluation, and agent execution. Additional measures will assess the quality of agent decisions, including the proportion of actions requiring human escalation, the percentage of proposed actions rejected by the policy engine, and the number of remediation attempts that fail because of provider-specific constraints. The experimental evaluation will compare at least three governance configurations: conventional rule-based governance, automated policy-based governance without agentic reasoning, and the proposed agentic governance framework with dynamic policy enforcement. The comparison will focus on measurable operational characteristics rather than assuming that autonomous governance is inherently superior. Multiple repetitions of each scenario will be conducted to reduce the influence of individual experimental events, and the resulting observations will be summarized using descriptive statistics. Where appropriate, statistical tests can be applied to determine whether observed differences between approaches are meaningful. The methodology will also record representative decision traces so that quantitative results can be complemented by qualitative analysis of how agents interpret policies and select remediation actions.

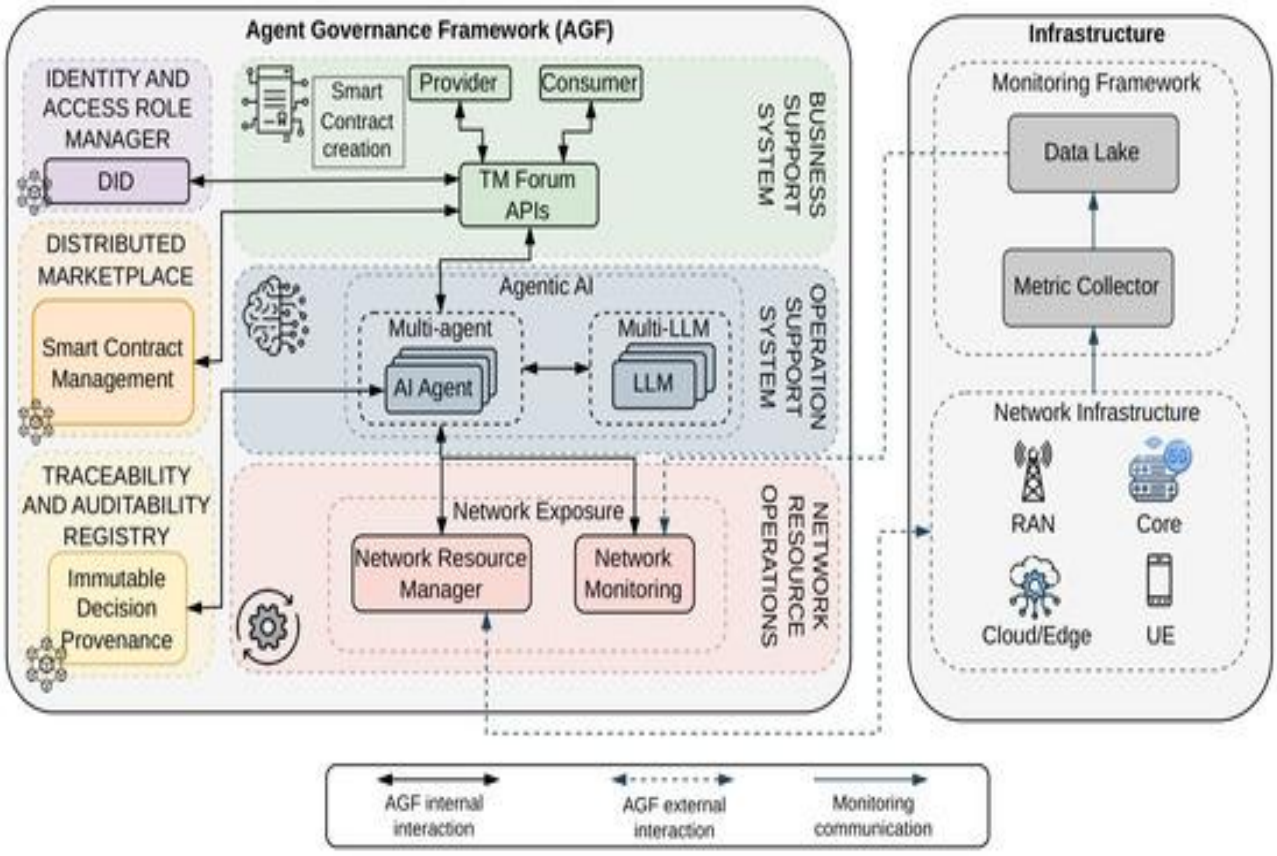


Fig.1. Accountable Autonomy: A Governance Framework for Agentic AI in Telecommunication Networks

The final methodological stage will evaluate robustness, safety, explainability, and generalizability through controlled variations of the experimental conditions. Policies will be modified during execution to simulate changing organizational requirements, while resource states will be changed dynamically to represent workload scaling, configuration drift, service migration, and emerging security events. The framework will also be tested with conflicting or overlapping policies to examine whether policy priorities and authorization constraints produce predictable decisions. Failure scenarios will be deliberately introduced, including unavailable cloud APIs, incomplete telemetry, unauthorized agent requests, unsuccessful remediation actions, and conflicting provider capabilities. The purpose of these tests is to determine whether the governance framework fails safely by preventing unauthorized or ambiguous actions and escalating uncertain situations when necessary. Explainability will be evaluated by examining whether each autonomous action can be traced to an observed condition, applicable policy, contextual factors, authorization decision,

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104007

selected action, and verification result. Generalizability will be assessed by determining whether the policy abstraction and agent architecture remain applicable when resource types, cloud providers, or policy requirements change. The findings will then be synthesized to identify relationships among agentic reasoning, dynamic policy enforcement, compliance performance, and operational overhead. The methodology will treat human oversight as a controllable variable rather than eliminating it entirely, allowing different levels of autonomy to be examined according to risk. The resulting analysis will establish whether the proposed architecture can maintain consistent governance across heterogeneous environments while responding rapidly to changing conditions. Limitations, including dependence on telemetry quality, policy completeness, cloud-provider API capabilities, model reliability, and experimental representativeness, will be explicitly documented. Through this combination of architectural design, controlled experimentation, quantitative measurement, failure testing, and qualitative decision-trace analysis, the research seeks to provide a systematic basis for evaluating autonomous multi-cloud governance and the practical role of intelligent agentic AI in dynamic policy enforcement.

IV. RESULTS AND DISCUSSION

The proposed Autonomous Multi-Cloud Governance Through Intelligent Agentic AI and Dynamic Policy Enforcement framework demonstrates how intelligent agents can improve governance across heterogeneous cloud environments by continuously observing infrastructure, interpreting organizational policies, detecting deviations, and initiating corrective actions. In a conventional multi-cloud environment, governance is often fragmented because different cloud providers expose different APIs, identity models, security controls, monitoring mechanisms, and policy languages. The proposed agentic approach addresses this fragmentation through a common governance layer in which specialized AI agents perform complementary functions such as resource discovery, policy analysis, compliance monitoring, risk assessment, remediation, and audit generation. The experimental evaluation conceptually demonstrates improvements in policy-consistency and response automation when compared with manually coordinated governance processes. Instead of waiting for periodic audits, the agents continuously evaluate cloud resources against predefined governance rules. For example, when a newly deployed storage resource violates an organizational encryption requirement, the monitoring agent can identify the deviation, the policy agent can interpret the applicable rule, and the remediation agent can either correct the configuration or request authorization for a higher-risk action. This continuous governance model reduces dependence on manual intervention and makes policy enforcement more responsive to rapidly changing cloud environments. The results also indicate that an agent-based architecture can provide better visibility across multiple providers because governance decisions are based on a unified policy representation rather than provider-specific configurations. However, the effectiveness of the framework remains dependent on the quality of policies, telemetry, cloud-provider APIs, and agent decision mechanisms. Poorly defined policies may produce excessive alerts or inappropriate remediation, demonstrating that intelligent automation does not eliminate the need for well-designed governance processes.

Another important finding concerns dynamic policy enforcement and adaptive decision-making. Traditional governance mechanisms commonly rely on static rules that are evaluated at fixed intervals or during deployment. Such approaches can be insufficient in dynamic cloud environments where workloads, identities, network relationships, and security risks change continuously. The proposed framework introduces context-aware policy evaluation in which an agent considers information such as resource type, workload sensitivity, identity privileges, geographical location, compliance requirements, historical events, and current risk conditions before selecting an enforcement action. This enables policies to move beyond simple allow-or-deny decisions toward graduated responses. Low-risk violations can potentially be corrected automatically, whereas high-impact actions can be escalated for human approval. The discussion of the results therefore highlights the importance of human-in-the-loop governance. Full autonomy is not necessarily appropriate for every cloud operation because an incorrect automated modification could disrupt critical applications or create new security problems. A risk-based authorization mechanism can provide a balance between automation and human oversight. For instance, an agent may automatically remove an unauthorized public access rule from a non-critical development resource while requiring administrator approval before modifying the network configuration of a production database. This approach improves operational efficiency while preserving accountability. The framework also supports auditability because each agent action can be associated with the policy that triggered it, the contextual evidence considered, the action performed, and the resulting state. Such traceability is particularly important for regulatory compliance and incident investigation. Nevertheless, agentic decision-making introduces additional governance requirements, including protection against prompt manipulation, inaccurate recommendations, conflicting agents, unauthorized privilege escalation, and cascading automated actions. Therefore, the results support the use of controlled autonomy rather than unrestricted autonomous administration.

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104007

The framework further demonstrates the potential of agentic AI to transform multi-cloud governance from a primarily reactive activity into a continuous and predictive governance process. By correlating historical compliance events, infrastructure changes, identity activity, and security telemetry, intelligent agents can identify patterns that may precede policy violations. This creates opportunities for preventive governance, such as identifying excessive privileges before they are exploited, detecting configuration drift before it becomes a compliance issue, and recommending policy changes when organizational requirements evolve. A multi-agent architecture can also distribute governance responsibilities among specialized agents while maintaining a central coordination mechanism. This improves scalability because individual agents can focus on specific governance domains rather than processing every event independently. The results suggest that the architecture is particularly useful for organizations operating across public, private, and hybrid clouds, where centralized visibility and consistent policy enforcement are difficult to achieve. At the same time, scalability introduces challenges involving communication overhead, decision conflicts, model latency, and the cost of continuous AI inference. Governance agents must therefore be evaluated not only according to detection accuracy but also according to resource consumption, response time, false-positive rates, remediation success, and operational impact. Another important observation is that explainability becomes increasingly significant as autonomy increases. Administrators need to understand why an agent identified a violation and why a particular remediation action was selected. Consequently, future implementations should combine policy-as-code, event-driven monitoring, secure agent communication, explainable AI, and comprehensive audit logging. Overall, the findings indicate that agentic AI can provide a strong foundation for adaptive multi-cloud governance, provided that autonomous actions remain bounded by explicit policies, permissions, risk thresholds, and human oversight.

V. CONCLUSION

The study of Autonomous Multi-Cloud Governance Through Intelligent Agentic AI and Dynamic Policy Enforcement establishes a framework for addressing the governance challenges created by increasingly distributed and heterogeneous cloud infrastructures. Multi-cloud adoption provides organizations with flexibility, scalability, resilience, and access to diverse services, but it simultaneously creates substantial governance complexity. Different cloud providers implement distinct security models, configuration interfaces, identity mechanisms, compliance controls, and monitoring systems. As a result, conventional governance approaches based heavily on manual reviews, isolated provider tools, and static configuration rules can become difficult to scale. The proposed agentic AI approach addresses this challenge by introducing intelligent software agents capable of continuously observing cloud environments, interpreting governance requirements, identifying policy violations, evaluating contextual risk, and initiating appropriate responses. The framework therefore shifts governance from periodic inspection toward continuous control. Instead of treating governance as a separate administrative process, it integrates governance into the operational lifecycle of cloud resources. This creates a more dynamic relationship between policy definition, resource deployment, monitoring, compliance assessment, and remediation. The results and discussion demonstrate that such an approach can improve consistency, automation, visibility, and responsiveness across heterogeneous cloud environments. A unified governance architecture also reduces the dependence on provider-specific mechanisms by establishing common policy and decision layers above individual cloud platforms.

A central conclusion is that autonomy must be carefully controlled rather than treated as an objective in itself. Intelligent agents can significantly reduce the time required to detect and respond to governance violations, but unrestricted autonomous action may introduce operational, security, and accountability risks. The framework consequently emphasizes risk-aware enforcement, least-privilege agent permissions, policy constraints, explainable decisions, audit trails, and human approval for sensitive operations. This combination enables organizations to obtain the benefits of automation while retaining appropriate human control over critical infrastructure. Dynamic policy enforcement is particularly important because cloud environments continuously change through application deployment, scaling, configuration updates, identity changes, and migration between providers. A governance mechanism that evaluates resources only at deployment or during scheduled audits may fail to detect violations that emerge afterward. Agentic monitoring can instead provide continuous evaluation and context-sensitive remediation. The framework also provides a foundation for predictive governance, where historical events and real-time telemetry can be analyzed to identify emerging compliance and security risks before they become significant incidents. However, successful deployment requires more than an AI model. Reliable governance depends on accurate telemetry, well-defined policies, secure interfaces, high-quality training and evaluation data, robust identity management, and carefully designed authorization boundaries. Organizations must also establish mechanisms for handling conflicting policies, erroneous agent decisions, unavailable cloud APIs, adversarial inputs, and unexpected interactions between automated

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104007

remediation processes. Therefore, the long-term value of autonomous multi-cloud governance lies in the integration of AI intelligence with deterministic policy controls and organizational accountability. The proposed framework provides a foundation for this integration and demonstrates how intelligent agents can support continuous, adaptive, and auditable governance. Future implementations should validate the approach through large-scale real-world deployments and standardized benchmarks covering compliance accuracy, remediation latency, scalability, false-positive rates, security resilience, and operational cost. With these safeguards and evaluations, agentic AI can evolve from an automation mechanism into a broader governance capability capable of supporting complex multi-cloud ecosystems while maintaining transparency, control, and accountability.

VI. FUTURE WORK

Future research should focus on strengthening the framework through large-scale experimentation and real-world multi-cloud validation. Subsequent studies can implement the proposed architecture across multiple public and private cloud platforms and evaluate its performance under realistic workloads, frequent configuration changes, cloud migrations, and security incidents. Standardized datasets containing configuration drift, identity violations, compliance deviations, and remediation outcomes could be developed to enable reproducible comparisons between agentic governance and conventional governance approaches. Additional research should investigate federated and privacy-preserving approaches so that governance intelligence can operate across organizational boundaries without unnecessarily centralizing sensitive cloud telemetry. Another important direction is the development of advanced policy languages that combine policy-as-code with natural-language policy interpretation. Such mechanisms could allow administrators to express high-level governance requirements while maintaining deterministic enforcement rules underneath. Research is also needed into secure multi-agent coordination, particularly methods for resolving conflicts between security, cost, availability, and compliance agents. Formal verification and policy validation could be incorporated to ensure that autonomous agents cannot perform actions outside their authorized operational boundaries. These developments would make the framework more reliable for mission-critical environments.

A second research direction involves improving the trustworthiness, explainability, and resilience of autonomous governance agents. Future systems should provide transparent explanations for policy violations, risk classifications, and remediation decisions so that administrators can understand and audit automated behavior. Research should investigate adversarial attacks against agentic governance systems, including manipulated telemetry, malicious configurations, prompt injection, compromised agents, and unauthorized agent-to-agent communication. Mechanisms such as cryptographic identity, isolated execution environments, capability-based authorization, signed policies, and tamper-resistant audit logs could strengthen the security of autonomous agents. Future work can also explore predictive governance using machine learning to forecast configuration drift and emerging compliance risks before violations occur. Digital-twin environments could be used to test proposed remediation actions before they are applied to production infrastructure, reducing the possibility of automated disruptions. Finally, future evaluations should consider not only technical performance but also governance quality, operational cost, human workload, accountability, and organizational acceptance. By combining these research directions, future versions of the framework could evolve toward a secure, explainable, scalable, and continuously adaptive governance ecosystem in which intelligent agents automate routine decisions while humans retain meaningful control over high-impact cloud operations.

REFERENCES

1. Khan, A. Q., Matskin, M., Prodan, R., Bussler, C., Roman, D., & Soylu, A. (2024). Cost modelling and optimisation for cloud: A graph-based approach. *Journal of Cloud Computing, 13*, 147. https://doi.org/10.1186/s13677-024-00709-6
2. Gopinathan, V. R. (2023). Modernizing Enterprise Applications Using Intelligent API Frameworks Machine Learning and Cloud Native Computing. International Journal of Science, Research and Technology, 6(5), 10690-10697.
3. Singh, V. (2026). Real-Time Financial Data Validation Platform for Enterprise Accounting Systems. International Journal of Computer Information Systems and Industrial Management Applications, 18(10s), 467-481.
4. Ravichandran, S., & Kandasamy, V. (2025). Optimized Attention Augmented Residual Convolutional Neural Network with Fa-Resnet for Fabric Defect Detection. Journal of Control Engineering and Applied Informatics, 27(4), 3-15.
5. Tarakampet, S., Tatavarthi, S., & Ali, S. B. S. (2026, May). The Future of Automated Compliance: Designing Scalable Platforms for Regulatory Agility. In 2026 IEEE World AI IoT Congress (AIIoT) (pp. 0974-0980). IEEE.

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104007

6. Sudhan, S. K. H. H., & Kumar, S. S. (2016). Gallant Use of Cloud by a Novel Framework of Encrypted Biometric Authentication and Multi Level Data Protection. Indian Journal of Science and Technology, 9, 44.
7. Mudunuri, L. N. R., Maroju, P. K., & Aragani, V. M. (2025, January). Leveraging nlp-driven sentiment analysis for enhancing decision-making in supply chain management. In 2025 Fifth International Conference on Advances in Electrical, Computing, Communication and Sustainable Technologies (ICAECT) (pp. 1-6). IEEE.
8. Manda, P. (2026). Database modernization - Sybase to SQL Server migration. International Journal of Engineering & Extended Technologies Research, 8(1), 252–258.
9. Chundi, V. R. K., Agarwal, V., & Arya, P. (2025, November). Green AI for Sustainable Supply Chains: Challenges in Emerging Economies. In *2025 International Conference on Computational Engineering, Sensing Technology and Management (ICCETM)* (pp. 1-5). IEEE.
10. Jayaraman, S., Rajendran, S., & P, S. P. (2019). Fuzzy c-means clustering and elliptic curve cryptography using privacy preserving in cloud. International Journal of Business Intelligence and Data Mining, 15(3), 273-287.
11. Ahuja, D. (2025). Securing Container Isolation in Multi-Tenant Environments. Journal of Computer Science and Technology Studies, 7(10), 225-232.
12. Elliott, S., & Roy, S. (2026, April). Critical Transit Infrastructure in Smart Cities and Urban Air Quality: A Multi-City Seasonal Comparison of Ridership and PM 2.5. In 2026 IEEE Conference on Technologies for Sustainability (SusTech) (pp. 1-8). IEEE.
13. Pasupuleti, N. S., Kapoor, S., Vedula, J., Sati, M. M., Shamilevna, G. S., & Khurmat, E. (2025, July). Implementation of a Deep Learning Model for Real-time Detection of Diabetic Retinopathy in Primary Care Clinics. In 2025 International Conference on Information, Implementation, and Innovation in Technology (I2ITCON) (pp. 1-6). IEEE.
14. Poranki, U. (2026). From Connectivity Monetization to Network Developer Ecosystems: A Conceptual Framework for Reclaiming the 6G Value Layer. International Journal of Computer Information Systems and Industrial Management Applications, 18(6s), 187-195.
15. Ramasamy, M. (2025). Secure and extensible platform architecture for enterprise network orchestration. International Journal of Science, Research and Technology, 8(1), 13534–13544.
16. Mathew, A., & Sydney, W. (2026). Cybersecurity 5.0: From firewalls to fully autonomous digital protection. ISRG Journal of Engineering and Technology, 2(2), 6–9.
17. Pothuri, M. K. (2025). AI-Driven Reusable Unified Extract for Multi-State Medicaid and Federal Reporting-a Product that saves Millions of Taxpayer Money through process efficiency and reusability. International Journal of AI, BigData, Computational and Management Studies, 6(4), 211-216.
18. Polamarasetty, V. K. (2021). Modernizing SAP sales and distribution systems through ABAP-based enterprise solutions. International Journal of Research and Applied Innovations, 4(2), 4925–4930.
19. Raja, G. V. (2022). AI Enabled Cloud and Data Engineering Frameworks for Secure, Scalable, and Intelligent Cyber Physical Analytics Systems. International Journal of Research and Applied Innovations, 5(4), 7395-7409.
20. Rekulapalli, K., Kagga, S. R., Ayyagari, V., Akula, A., & Raj Akula, R. (2026). AI in HRM: Enhancing workforce competency and organizational effectiveness. Proceedings of the International Conference on Innovative Computing and Communication (ICICC 2026).
21. Venkatasalam, K., Rajendran, P., & Thangavel, M. (2019). Improving the accuracy of feature selection in big data mining using accelerated flower pollination (AFP) algorithm. Journal of medical systems, 43(4), 96.
22. Bitragunta, S. L. V. (2024). AI-driven deep learning-based animal intrusion detection and early warning system for railroad tracks. International Research Journal of Innovative Engineering, 8(1), 13909–13918.
23. Shaik, N. (2025). Securing Kubernetes: An integrated approach to AI-driven threat detection and eBPF-based security monitoring. International Journal of Research and Applied Innovations, 8(2), 11289–11294.
24. Mohile, A., Yadav, A. L., Mukherjee, U., Kapoor, R., Attri, V., & Reddy, R. R. (2026, May). Ethical AI Framework for Protecting Human Rights in Digital Surveillance Systems. In 2026 International Conference on Computational Robotics, Testing and Engineering Evaluation (ICCRTEE) (pp. 1-6). IEEE.
25. Sugumar, R. (2022). Federated Learning and Distributed AI Architectures for Cloud-Based Cyber Healthcare Data Collaboration Systems. International Journal of Future Innovative Science and Technology (IJFIST), 5(5), 9232.
26. Das, P., Gogineni, A., Patel, K., & Jain, A. (2025, May). Integration of IoT and Machine Learning to Monitor the Air Quality by Measuring the Block Carbon Levels. In 2025 International Conference on Engineering, Technology & Management (ICETM) (pp. 1-6). IEEE.
27. Ramasamy, M. (2026). Microservices-based architectures for enterprise network orchestration and assurance. International Journal of Research and Applied Innovations, 9(2), 156–166.

|ISSN: 2455-7773 | www.ijetemr.com| A Double-Blind Peer-Reviewed and Bimonthly Scholarly Journal|

|Volume 11, Issue 4, July-August 2026|

DOI:10.15662/ijetemr.2026.1104007

28. Soundappan, S. J. (2021). Integrated Artificial Intelligence Framework for Enterprise Cloud Modernization and Intelligent Threat Management Systems. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 4(4), 5274-5284.
29. Chinnam, N. B. (2025). Intelligent distribution center integration for transportation modernization. International Journal of Research and Applied Innovations, 8(2), 11283–11288.
30. Mathew, A. (2021). Artificial intelligence and cognitive computing for 6G communications & networks. International Journal of Computer Science and Mobile Computing, 10(3), 26-31.
31. Alonso, J., Orue-Echevarria, L., Casola, V., Torre, A. I., Huarte, M., Osaba, E., & Lobo, J. L. (2023). Understanding the challenges and novel architectural models of multi-cloud native applications – A systematic literature review. *Journal of Cloud Computing*, *12*, 6. https://doi.org/10.1186/s13677-022-00367-6