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 10, Issue 6, November-December 2025|

DOI:10.15662/ijetemr.2025.1006011

Explainable Machine Intelligence Frameworks for Cloud Security Posture Management in Regulated Enterprise Environments

Venkata Sri Manoj Bonam

Data Engineer, New York Life Insurance, New York, USA

ABSTRACT: Explainable Machine Intelligence (EMI) is emerging as an important approach for improving Cloud Security Posture Management (CSPM) in regulated enterprise environments where security decisions must be transparent, auditable, and aligned with organizational and regulatory requirements. Modern cloud infrastructures contain dynamic workloads, distributed identities, complex configurations, containerized applications, application programming interfaces, and continuously changing security controls. Traditional CSPM approaches primarily depend on predefined rules, configuration checks, and compliance mappings, which can generate large volumes of alerts without adequately explaining their context, significance, or remediation priorities. Machine intelligence can enhance CSPM through automated anomaly detection, risk identification, configuration analysis, behavioral monitoring, and predictive security assessment. However, opaque machine-learning models can create challenges for auditors, security teams, and regulated organizations when decisions cannot be adequately interpreted or justified. This paper proposes a conceptual explainable machine intelligence framework that integrates machine-learning analytics, cloud telemetry, security posture assessment, regulatory control mapping, explainability mechanisms, and human oversight. The framework emphasizes interpretable risk reasoning, traceable evidence, confidence information, policy alignment, and auditable remediation recommendations. It also considers challenges involving data privacy, model drift, adversarial manipulation, false positives, explainability limitations, and accountability. The proposed framework provides a foundation for developing transparent, adaptive, and governance-oriented CSPM architectures capable of supporting security operations while preserving regulatory accountability and human decision-making.

KEYWORDS: Explainable Artificial Intelligence, Machine Intelligence, Cloud Security, Cloud Security Posture Management, CSPM, Enterprise Security, Regulatory Compliance, Cybersecurity, Risk Management, Machine Learning, Security Governance, Cloud Computing, Auditability, Threat Detection

I. INTRODUCTION

Cloud adoption has fundamentally changed the security landscape of regulated enterprises by distributing infrastructure, applications, identities, data, and security responsibilities across increasingly complex computing environments. Organizations may operate across multiple cloud providers, hybrid infrastructures, private cloud environments, containers, serverless workloads, and software-as-a-service platforms. Such environments can change rapidly as resources are created, modified, scaled, and removed. Consequently, maintaining an accurate security posture requires continuous visibility into configurations, identities, vulnerabilities, network exposure, data protection, and compliance controls. Cloud Security Posture Management (CSPM) has emerged as an important approach for identifying misconfigurations, security weaknesses, compliance deviations, and potentially risky cloud resources. However, conventional rule-based CSPM systems can produce large numbers of findings without adequately explaining why a finding matters, how different risks are related, or why a particular remediation action should receive attention. Machine intelligence can extend CSPM by analyzing large volumes of security telemetry and identifying patterns that may not be captured by static rules.

The increasing use of machine-learning models in security operations introduces a corresponding requirement for explainability. Security analysts and compliance teams need to understand the evidence supporting a machine-generated risk assessment, particularly when the affected systems contain sensitive information or support regulated business processes. Explainability can help connect technical findings with business context, regulatory controls, security policies, and remediation procedures. For example, instead of simply identifying an insecure cloud configuration, an explainable system could indicate the relevant configuration, the observed evidence, the applicable policy requirement, the factors contributing to the risk assessment, and the recommended corrective action. Such capabilities can improve investigation and audit processes while allowing security professionals to challenge or validate

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

|Volume 10, Issue 6, November-December 2025|

DOI:10.15662/ijetemr.2025.1006011

automated recommendations. However, explanations must themselves be accurate, understandable, secure, and traceable. An explanation that does not faithfully represent the model's reasoning may create a misleading sense of assurance. This study therefore focuses on an integrated framework in which machine intelligence is combined with explainability, cloud security posture assessment, regulatory mapping, evidence management, and human oversight. The objective is to establish a conceptual foundation for transparent and accountable CSPM in regulated enterprise environments.

II. LITERATURE REVIEW

Research in cloud security has established that cloud environments introduce distinctive challenges related to shared responsibility, identity management, resource configuration, network exposure, workload isolation, data protection, and visibility across distributed infrastructure. CSPM approaches traditionally address these challenges through continuous configuration assessment, policy enforcement, compliance checking, vulnerability identification, and security monitoring. Rule-based systems are effective for deterministic requirements because they can directly evaluate whether a configuration satisfies a specified condition. However, complex enterprise environments generate relationships among resources, identities, vulnerabilities, workloads, and business processes that may not be fully represented through isolated rules. Machine-learning approaches provide an opportunity to analyze these relationships and identify anomalies, unusual behavior, risk patterns, and potentially significant combinations of security conditions. The literature on intelligent cybersecurity consequently emphasizes the potential of machine learning to supplement conventional detection mechanisms while recognizing challenges involving training data, false positives, model drift, adversarial attacks, and generalization across changing environments.

Explainable Artificial Intelligence (XAI) research provides an additional foundation for the proposed framework. Explainability encompasses methods that help users understand why an AI system produced a particular output, which features influenced a prediction, and what evidence is associated with a decision. Techniques can include feature attribution, interpretable models, counterfactual explanations, rule extraction, local explanations, and model-agnostic interpretation approaches. In cybersecurity, explainability can be particularly valuable because analysts often need to validate automated alerts before taking corrective action. An explanation may identify contributing variables such as unusual access patterns, publicly exposed resources, excessive privileges, configuration deviations, vulnerability severity, or relationships among cloud assets. Nevertheless, explainability is not synonymous with accuracy or fairness. An explanation may be understandable while failing to faithfully describe the underlying model behavior. Therefore, enterprise security systems require mechanisms for validating explanation fidelity and preserving evidence supporting machine-generated assessments.

Regulatory and governance literature further emphasizes accountability, traceability, risk management, access control, data protection, and auditability. Regulated enterprises may need to demonstrate that security controls are appropriately implemented and that significant security activities can be documented and reviewed. This requirement creates an important relationship between CSPM and explainable machine intelligence. A machine-generated security finding can provide greater governance value when it is connected to a specific asset, configuration state, security policy, regulatory control, evidence record, risk rationale, and remediation history. Frameworks for responsible AI similarly emphasize human oversight, transparency, security, documentation, and lifecycle management. These principles suggest that explainable CSPM should not simply add a natural-language explanation to an automated alert. Instead, explainability should be integrated throughout the security assessment lifecycle, from data collection and feature generation to risk scoring, control mapping, remediation recommendations, and audit logging. The proposed research therefore combines insights from cloud security, CSPM, machine learning, explainable AI, and security governance to conceptualize an architecture capable of producing technically useful and auditable security assessments.

III. RESEARCH METHODOLOGY

This study adopts a qualitative conceptual research methodology to develop an explainable machine intelligence framework for Cloud Security Posture Management in regulated enterprise environments. The methodology is designed to integrate concepts from cloud cybersecurity, machine learning, explainable artificial intelligence, security governance, and regulatory compliance into a unified analytical framework. The research begins by identifying the principal components of the problem domain, including cloud assets, security configurations, identities, vulnerabilities, network relationships, telemetry, compliance requirements, machine-learning models, explanation mechanisms, security analysts, and governance processes. These components are organized into four conceptual dimensions: cloud

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

|Volume 10, Issue 6, November-December 2025|

DOI:10.15662/ijetemr.2025.1006011

security data, intelligent security analytics, explainability and evidence, and governance and human oversight. Cloud security data includes configuration records, identity information, vulnerability information, network relationships, activity logs, threat intelligence, asset inventories, and policy states. Intelligent security analytics encompasses anomaly detection, risk classification, behavioral analysis, misconfiguration detection, attack-path analysis, and predictive assessment. Explainability encompasses feature attribution, evidence presentation, confidence information, counterfactual reasoning, rule extraction, and natural-language explanations. Governance and oversight include regulatory control mapping, audit trails, approval mechanisms, access control, policy enforcement, model monitoring, and human review. The framework is designed around the principle that machine intelligence should support security decision-making without removing accountability from authorized security professionals and governance functions.

The second methodological stage involves thematic analysis of relevant research and established cybersecurity practices. Literature concerning CSPM is examined to identify conventional posture-management functions such as continuous configuration assessment, compliance monitoring, vulnerability identification, asset discovery, and remediation management. Research concerning machine learning in cybersecurity is analyzed to identify appropriate applications for anomaly detection, classification, behavioral modeling, and risk analysis, while also documenting limitations involving false positives, incomplete training data, model drift, adversarial manipulation, and changing cloud architectures. XAI literature is reviewed to categorize explanation mechanisms according to their interpretability, technical requirements, and suitability for security operations. Regulatory and governance literature is examined to identify requirements related to evidence, accountability, privacy, auditability, access control, and risk management. The analysis distinguishes between descriptive evidence and machine-generated interpretation. For example, the existence of a public network exposure can be established as an observable configuration fact, whereas the degree of business risk associated with that exposure may require contextual analysis. This distinction is important because an explainable framework should clearly separate observed evidence, model inference, policy interpretation, and recommended remediation. The thematic analysis consequently establishes a basis for designing a system in which explanations are grounded in traceable security evidence rather than generated solely as unrestricted natural-language narratives.

The third stage develops the proposed framework as a multilayer architecture. The first layer is a cloud telemetry and asset-observation layer that collects configuration states, identity relationships, network information, vulnerability data, logs, resource metadata, and security events from enterprise cloud environments. The second layer is a normalization and knowledge layer that establishes consistent representations of cloud resources, relationships, controls, and security events across heterogeneous platforms. This layer can maintain an asset graph connecting resources, identities, applications, vulnerabilities, permissions, network paths, and regulatory controls. The third layer is the machine intelligence layer, where supervised, unsupervised, graph-based, or hybrid analytical models identify anomalies, detect risk patterns, classify posture conditions, and estimate potential security significance. The fourth layer is the explainability layer, which associates each machine-generated finding with relevant input features, observed evidence, confidence information, contributing conditions, and potentially applicable counterfactuals. The fifth layer is the governance layer, which maps findings to organizational policies and regulatory controls, records evidence, maintains audit trails, applies authorization requirements, and manages remediation workflows. Finally, the human oversight layer allows security analysts, compliance specialists, and authorized administrators to review findings, challenge interpretations, approve remediation actions, and provide feedback. This architecture follows a human-in-the-loop model in which machine intelligence performs large-scale analysis while humans retain responsibility for consequential security decisions.

The final methodological stage evaluates the conceptual framework using analytical criteria appropriate for regulated enterprise security environments. The first criterion is detection capability, assessing whether the framework can identify relevant configuration deviations, behavioral anomalies, and interconnected risk conditions. The second is explanation quality, examining whether generated explanations are understandable, evidence-based, sufficiently faithful to model behavior, and useful for security investigation. The third is auditability, assessing whether security findings can be traced from the observed cloud state through model processing, explanation generation, policy mapping, and remediation activity. The fourth is regulatory alignment, examining whether technical findings can be associated with applicable organizational and regulatory requirements without treating machine-generated interpretations as authoritative legal conclusions. The fifth criterion is security and privacy, focusing on protection of telemetry, sensitive enterprise information, model interfaces, and explanation outputs. The sixth is operational usability, considering whether explanations help analysts investigate findings without creating additional alert or documentation burdens. The seventh is resilience, addressing model drift, adversarial manipulation, changing cloud architectures, and incomplete or

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

|Volume 10, Issue 6, November-December 2025|

DOI:10.15662/ijetemr.2025.1006011

noisy telemetry. These criteria provide a structured basis for subsequent empirical validation through enterprise case studies, controlled experiments, benchmark datasets, analyst evaluations, and longitudinal operational studies. The methodology does not assume that explainability automatically produces more accurate security decisions; instead, it treats explanation as one component of a broader assurance architecture involving evidence, validation, governance, and human judgment. The resulting framework positions explainable machine intelligence as an analytical capability within CSPM rather than an autonomous authority. Its central proposition is that regulated enterprises can obtain greater transparency from intelligent security automation when machine-generated findings are connected to observable cloud evidence, interpretable analytical factors, explicit security policies, regulatory mappings, and auditable human decisions. This conceptual foundation can support future empirical research measuring detection performance, explanation fidelity, analyst effectiveness, compliance traceability, remediation efficiency, and organizational trust in AI-assisted cloud security management.

IV. RESULTS AND DISCUSSION

The results indicate that an Explainable Machine Intelligence (EMI) framework can substantially improve the contextual analysis and management of cloud security posture in regulated enterprise environments by combining machine-learning analytics, cloud configuration intelligence, compliance policies, asset context, and human-readable explanations. Traditional Cloud Security Posture Management (CSPM) approaches are generally effective at detecting known configuration violations, insecure permissions, exposed services, missing security controls, and compliance deviations, but they can produce large volumes of findings without adequately explaining their business significance. The proposed framework addresses this limitation by establishing a multi-layer architecture consisting of cloud-data acquisition, security-context enrichment, machine-intelligence analysis, explainability, governance, and remediation. The first result is improved visibility across heterogeneous cloud environments. The framework can collect configuration and security information from virtual machines, containers, Kubernetes clusters, storage services, databases, serverless resources, networks, identity and access management systems, and security-monitoring platforms. These technical observations can then be enriched with information such as asset ownership, business criticality, data sensitivity, network exposure, identity privileges, vulnerability information, and applicable organizational policies. This contextualization changes the nature of CSPM from simple configuration checking to risk-aware posture analysis. For example, an excessive permission assigned to an isolated development resource can be distinguished from an equivalent permission associated with a production database containing regulated information. Machine intelligence can correlate these attributes and identify relationships that may not be visible through independent rule checks. Graph-based representations are particularly useful because they can model relationships among users, roles, workloads, network connections, datasets, vulnerabilities, and policies. The resulting security graph provides a foundation for identifying potentially significant attack paths and combinations of weaknesses. At the same time, the results demonstrate that machine-learning outputs must be interpreted cautiously because model decisions depend on data quality, feature selection, training assumptions, and organizational risk definitions. Consequently, the framework should preserve source evidence and expose the factors that contribute to each finding rather than presenting an unexplained numerical risk score. Explainability therefore becomes an architectural component rather than merely a visualization feature. Local explanations can describe why a particular cloud resource was classified as risky, global explanations can identify recurring patterns in model behaviour, and counterfactual explanations can indicate which configuration changes could alter a particular security decision. This combination provides security analysts with greater context for investigation while allowing governance teams to trace findings back to identifiable technical evidence. Overall, the results suggest that the integration of machine intelligence and explainability can increase the usefulness of CSPM outputs by reducing the distance between automated detection and human security decision-making.

A second major result concerns the framework's contribution to regulatory compliance, auditability, and intelligent governance. Regulated enterprises operate under organizational security policies and external requirements involving data protection, access management, retention, encryption, logging, incident management, and information security controls. Conventional compliance monitoring frequently relies on predefined rules and periodic assessments. Although these mechanisms remain necessary, they can become difficult to maintain in rapidly changing cloud environments. The proposed framework introduces continuous governance assistance by connecting cloud-state information with machine-readable policies and contextual evidence. When a new resource is deployed, the framework can evaluate its configuration against applicable controls; when a configuration changes, it can reassess the resource's security posture; and when a potential policy violation is identified, the system can explain the affected control, resource, evidence, and potential remediation. Natural-language interfaces can also allow authorized users to query security posture information

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

|Volume 10, Issue 6, November-December 2025|

DOI:10.15662/ijetemr.2025.1006011

without manually navigating multiple cloud-management interfaces. For example, a compliance analyst could request an explanation of which production resources have externally accessible interfaces or which resources lack a required security control. The framework can translate complex relationships into role-specific explanations for security engineers, auditors, compliance officers, and business stakeholders. This capability can improve communication because each group can receive an explanation appropriate to its responsibilities while retaining access to the same underlying evidence. Another important result is the development of an auditable evidence chain. Each machine-generated finding can be associated with the source configuration, timestamp, relevant policy, model or analytical method, explanation, recommended remediation, human approval, remediation outcome, and subsequent verification. Such traceability is particularly valuable in regulated environments where organizations may need to demonstrate how security controls were monitored and how exceptions were handled. However, the results also identify an important limitation: generative models can produce explanations that are linguistically convincing but technically inaccurate. An explanation should therefore not be treated as reliable merely because it is understandable. The proposed architecture addresses this problem by grounding explanations in authoritative enterprise sources such as cloud configuration records, approved policies, security-control definitions, asset inventories, and verified telemetry. Retrieval mechanisms can restrict the information available to the model according to user authorization, while deterministic compliance rules remain authoritative for formal control evaluation. This hybrid approach establishes a division of responsibility in which machine intelligence supports interpretation, correlation, prioritization, and explanation, whereas deterministic controls establish enforceable security and compliance conditions. The results further indicate that this architecture can improve audit readiness because security teams can obtain structured evidence continuously rather than reconstructing historical cloud states manually. Nevertheless, continuous monitoring introduces additional requirements for secure logging, data retention, privacy protection, access control, and model governance. Explainable CSPM should therefore be integrated into the enterprise governance architecture rather than deployed as an isolated artificial-intelligence service.

The third major result relates to risk prioritization, remediation, operational efficiency, and human-machine collaboration. Large enterprise cloud environments can generate substantial numbers of security findings, and treating every finding equally can overwhelm security operations teams. The framework uses contextual machine intelligence to correlate configuration weaknesses with asset criticality, external exposure, privilege levels, data sensitivity, vulnerability information, and relationships among cloud resources. This enables security teams to investigate related findings as interconnected security conditions rather than as independent alerts. For instance, a publicly reachable workload, an excessive identity privilege, and an unpatched component may individually represent separate findings, but their combination may create a more meaningful security scenario. Graph-based analysis can identify such relationships and provide an interpretable representation of the contributing resources and controls. Explainability then enables analysts to understand why the system associated those resources and what evidence supports the resulting prioritization. This can help reduce alert fatigue and focus investigation on findings with stronger contextual relevance. The framework can additionally generate remediation recommendations, including configuration changes, access-policy modifications, infrastructure-as-code updates, or additional monitoring requirements. However, automated remediation requires careful controls because security changes can affect application availability and legitimate business processes. The results therefore support a staged remediation architecture consisting of detection, explanation, recommendation, approval, execution, and verification. Low-impact and previously authorized corrections may be automated according to predefined policies, whereas changes affecting critical production resources should require appropriate human authorization. After remediation, the system should independently verify whether the intended security state was achieved and whether the change introduced secondary effects. Evaluation of the framework should consequently extend beyond detection accuracy. Relevant measures include precision, recall, false-positive and false-negative rates, remediation time, policy-coverage percentage, explanation fidelity, explanation completeness, analyst comprehension, investigation time, and remediation accuracy. Human-centred testing is especially important because the objective of explainability is not simply to generate longer descriptions but to improve the ability of security professionals to understand and validate machine-generated findings. The discussion therefore indicates that explainability should be measured according to whether users can identify the evidence behind a finding, understand the principal contributing factors, recognize uncertainty, and determine whether a recommended action is appropriate. The results also highlight the importance of model governance. Changes in cloud architectures, policies, threat patterns, and enterprise applications can cause model performance to change over time. Continuous monitoring of model behaviour, periodic validation, version control, and independent review are consequently required. Overall, the findings support an architecture in which machine intelligence provides scalable analytical capabilities while deterministic controls, evidence preservation, and human oversight maintain accountability. Such an approach can make CSPM more

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

|Volume 10, Issue 6, November-December 2025|

DOI:10.15662/ijetemr.2025.1006011

contextual, transparent, and operationally useful without treating machine-generated decisions as inherently authoritative.

V. CONCLUSION

The findings demonstrate that Explainable Machine Intelligence can provide a structured foundation for improving Cloud Security Posture Management in regulated enterprise environments. The proposed framework extends conventional CSPM beyond static configuration checking by combining continuous cloud-data collection, contextual security enrichment, machine-learning analysis, graph-based relationship modelling, explainability mechanisms, compliance intelligence, and controlled remediation. This combination allows security teams to move from isolated alerts toward contextual understanding of security conditions. The framework can identify relationships between cloud resources, identities, vulnerabilities, network exposure, data sensitivity, and organizational policies, thereby providing a broader representation of enterprise security posture. Its explainability component is particularly important because machine-generated risk classifications without supporting evidence can be difficult to validate and may not satisfy the transparency requirements of regulated organizations. By associating findings with source configuration, relevant policies, analytical factors, and remediation evidence, the framework can make machine-assisted security operations more traceable. Local explanations can clarify why a particular resource was identified as problematic, global explanations can expose broader patterns in model behaviour, and counterfactual explanations can demonstrate how specific configuration changes could affect a security decision. These capabilities support communication among security engineers, compliance professionals, auditors, and business stakeholders. The study also establishes that explainability must be connected to trustworthy evidence. Generative models can create plausible but inaccurate explanations, making evidence grounding, retrieval restrictions, deterministic policy enforcement, and human validation essential. Consequently, the framework should not replace established security controls or organizational governance mechanisms. Instead, machine intelligence should operate as an analytical and decision-support layer above authoritative cloud configurations, security policies, access controls, audit systems, and compliance rules. This distinction is critical for regulated environments because accountability cannot be delegated solely to a probabilistic model. The framework's value therefore arises from combining machine-scale analysis with mechanisms that preserve human responsibility, traceability, and control.

The study further concludes that explainable CSPM can contribute to more efficient security operations when risk analysis, governance, and remediation are integrated into a common architecture. Contextual prioritization can help security teams distinguish between isolated configuration weaknesses and combinations of conditions that require additional investigation. Explainable recommendations can provide analysts with the evidence and reasoning needed to evaluate proposed remediation actions, while controlled automation can reduce repetitive operational work. Nevertheless, enterprise deployment requires careful attention to model limitations, data quality, privacy, security, inference costs, and operational impact. A security model trained or configured using incomplete cloud information may produce unreliable conclusions, while changes in infrastructure or organizational policy may reduce performance over time. Therefore, continuous validation, model monitoring, version management, and human review should form part of the framework's operating model. The evaluation of explainable CSPM should also use measurable indicators such as detection accuracy, false-positive rates, policy coverage, remediation time, explanation fidelity, analyst comprehension, and audit-trail completeness. These measures provide a more comprehensive assessment than simply counting the number of findings detected. The resulting framework can be viewed as a governance-aware intelligence layer that connects cloud telemetry with security knowledge and human decision-making. Its purpose is not to eliminate security professionals but to provide them with scalable analytical capabilities, contextual evidence, and understandable explanations. When implemented with strong access controls, deterministic compliance rules, secure data handling, auditable workflows, and human approval for sensitive actions, explainable machine intelligence can support continuous and transparent cloud security posture management. The broader significance of this approach is that regulated enterprises can pursue greater automation without abandoning accountability. Future enterprise cloud-security architectures can therefore combine intelligent analytics with explicit governance and verifiable evidence, creating security operations that are increasingly adaptive while remaining understandable, auditable, and controllable.

VI. FUTURE WORK

Future work should focus on developing and empirically validating explainable CSPM frameworks across multi-cloud and hybrid-cloud environments. Research can investigate how machine intelligence can consistently interpret security information from different cloud providers, container platforms, Kubernetes environments, identity systems,

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

|Volume 10, Issue 6, November-December 2025|

DOI:10.15662/ijetemr.2025.1006011

infrastructure-as-code repositories, and legacy enterprise applications. Security knowledge graphs represent another promising direction because they can model relationships between identities, permissions, workloads, vulnerabilities, network paths, data assets, and compliance controls. Combining graph analytics with machine learning may enable more comprehensive identification of compound security conditions while retaining an interpretable representation of the relationships involved. Future studies should also establish standardized evaluation datasets containing realistic cloud configurations, security misconfigurations, compliance requirements, and attack-path scenarios. Such datasets could support reproducible comparison of detection precision, recall, false-positive rates, explanation fidelity, remediation accuracy, and computational cost. Research should additionally investigate the effectiveness of different explainability techniques, including feature-based explanations, counterfactual reasoning, graph explanations, rule-based explanations, and retrieval-grounded natural-language explanations. Human-centred experiments involving security engineers and compliance professionals will be important for determining whether explanations genuinely improve investigation and decision-making rather than simply increasing the amount of information presented to users.

A second research direction involves trustworthy autonomous remediation and continuous regulatory governance. Future frameworks could use specialized AI agents to monitor posture, investigate correlated findings, generate remediation plans, and verify corrective actions while remaining constrained by policy-as-code, role-based permissions, approval workflows, and immutable audit records. Research should determine which categories of remediation can safely be automated and which require human authorization based on resource criticality and potential operational impact. Privacy-preserving machine learning and confidential-computing techniques could also be investigated to protect sensitive enterprise security information during model inference and training. Another important area is model lifecycle governance. Long-term research should examine how model explanations, risk classifications, and detection performance change as cloud architectures, security policies, vulnerabilities, and organizational requirements evolve. Continuous drift detection, model versioning, independent validation, and automated regression testing could help maintain reliability. Finally, future studies should evaluate the economic dimension of explainable CSPM by measuring cloud-computing costs, inference overhead, analyst time saved, remediation efficiency, and compliance-monitoring effort. These investigations can contribute to practical reference architectures and implementation standards for regulated enterprises seeking to integrate machine intelligence into cloud-security operations while preserving transparency, accountability, privacy, and human oversight.

REFERENCES

1. Devineni, A., Pochincharla, S., & Muppalla, L. K. (2022). A telemetry-fusion framework for proactive incident detection in multi-tenant financial cloud platforms. International Journal of Advances in Signal and Image Sciences, 13–18.
2. Bitragunta, S. L. V. (2024). Intelligent electric vehicle charging hubs AI-based demand forecasting and smart energy management. International Journal of Emerging Trends in Engineering and Management Research, 9(2), 15383–15393.
3. Roy, S., & Sen, A. (2021, January). A self-updating K-contingency list for smart grid system. In 2021 IEEE 11th Annual Computing and Communication Workshop and Conference (CCWC) (pp. 0341-0346). IEEE.
4. Shaik, N. (2025). Artificial intelligence in healthcare economics: Transforming cost efficiency, member experience, and preventative care. International Journal of Research Publications in Engineering, Technology and Management, 8(2), 11850–11854.
5. Neela, S. (2023). Performance Benchmarking Methodologies for Real-Time Integration Middleware Platforms. American International Journal of Computer Science and Technology, 5(5), 39-46.
6. Bellundagi, M. (2022). Performance Optimization Techniques for Enterprise Java Applications Using Middleware and Messaging Systems. International Journal of Computer Technology and Electronics Communication, 5(3), 5158-5168.
7. Mathew, A., & Alex, H. (2025). Federated Learning for Secure Genomic Research: Privacy-Preserving AI Solutions for Precision Medicine. Science and Technology: Developments and Applications, 9, 36-43.
8. Yadavalli, V. R. (2022). Continuous trust scoring for cross-layer access governance in S/4HANA, Fiori, and SAP BTP. International Journal of Engineering & Extended Technologies Research, 4(1), 4347–4357.
9. Selvarajan, K. (2023). Optimizing large-scale data lakehouse architectures for performance and cost efficiency. International Journal of Research and Applied Innovations, 6(6), 10093–10101.
10. Pothuri, M. K. (2025). AI-powered governance with modern data engineering for regulatory excellence in healthcare and fintech. International Journal for Multidisciplinary Research, 7(4). https://doi.org/10.36948/ijfmr.2025.v07i04.52676

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

|Volume 10, Issue 6, November-December 2025|

DOI:10.15662/ijetemr.2025.1006011

11. 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.
12. 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.
13. Tarigoppula, S. (2022). Zero-trust security architecture for hybrid enterprise cloud environments using AWS native services. International Research Journal of Innovative Engineering, 6(6), 11433–11444.
14. Punithavathi, R., Selvi, R. T., Latha, R., Kadiravan, G., Srikanth, V., & Shukla, N. K. (2022). Robust node localization with intrusion detection for wireless sensor networks. Intelligent Automation and Soft Computing, 33(1), 143-156.
15. 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.
16. Chinnam, N. B. (2023). Modernizing retail fulfillment operations through automated load execution and microservices. International Journal of Computer Technology and Electronics Communication, 6(4), 7367–7371.
17. 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.
18. Ahuja, D. (2025). DevOps and Ethical AI: Ensuring Responsible Deployment. Journal Of Multidisciplinary, 5(6), 1-14.
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. 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.
21. Karakondu, M., Jambagi, G., & Tatavarthi, S. (2025). Optimising data loss prevention (DLP) strategies in cloud-native financial platforms.
22. Mathew, A. (2021). Artificial intelligence and cognitive computing for 6G communications & networks. International Journal of Computer Science and Mobile Computing, 10(3), 26-31.
23. Polamarasetty, V. K. (2022). Enterprise SAP application modernization for post-acquisition business transformation. International Journal of Science, Research and Technology (IJSRAT), 5(3), 7777–7783. https://www.ijsrat.com/index.php/ijsrat/issue/view/23
24. Ramasamy, M. (2023). Cloud-native control plane design for modern network automation systems. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 6(4), 9082–9091.
25. 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.
26. Srikanth, V., Walia, R., Augustine, P. J., Simla, J., & Jegajothi, B. (2022, March). Chaotic Whale Optimization based Node Localization Protocol for Wireless Sensor Networks Enabled Indoor Communication. In 2022 International Conference on Electronics and Renewable Systems (ICEARS) (pp. 702-707). IEEE.
27. Chinnam, N. B. (2024). Cross-network scheduling optimization for unified retail distribution networks. International Journal of Engineering & Extended Technologies Research (IJEETR), 6(4), 8183–8188.
28. 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.
29. Rahul Reddy Bandhela, RamMohan Reddy Kundavaram. (2023). A Comparative Study on Neural Network Architectures for Image Recognition Applications. Journal of Computational Analysis and Applications (JoCAAA), 31(1), 1334–1342. Retrieved from https://www.eudoxuspress.com/index.php/pub/article/view/3566