iJRASET
International Journal For Research in Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** VII **Month of publication:** July 2026

DOI: https://doi.org/10.22214/ijraset.2026.84086

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VII July 2026- Available at www.ijraset.com

An AI-Driven Adaptive Risk Assessment Framework for Secure and Resilient Cloud Infrastructure

Gandikota Nikhileshwari[1], N. Shiva Lingam[2]
[1, 2]*Department of Computer Science and Engineering, Holy Mary Institute of Technology & Science, Hyderabad, India*

Abstract: ***Cloud infrastructure is increasingly used to host academic, enterprise, and public-sector services, but its dynamic and distributed nature makes traditional static risk assessment insufficient. Static rule-based scoring is often unable to com-bine vulnerability severity, threat likelihood, workload impact, anomalous behavior, and recent operational context into a single actionable risk view. This paper proposes an AI-driven adap-tive risk assessment framework for secure and resilient cloud infrastructure. The framework integrates cloud log collection, preprocessing, Isolation Forest-based anomaly detection, Random Forest-based risk classification, weighted adaptive risk scoring, alert generation, mitigation recommendation, and dashboard-based visualization. The proposed risk engine computes a nor-malized score using vulnerability score, threat probability, impact score, and anomaly risk. The system was implemented as a defensive Flask-based prototype using simulated cloud infrastruc-ture logs. Experimental evaluation on 600 simulated cloud events produced an average risk score of 39.6, identified 108 high-risk and 13 critical-risk events, generated 746 explainable alerts, and produced 1048 defensive recommendations. The proposed hybrid model achieved 95.20% accuracy and a 93.75% F1-score in risk-level classification, outperforming a static rule-based baseline. The results demonstrate that adaptive AI-assisted risk assessment can improve cloud security monitoring, decision support, and resilience planning.***

Index Terms: ***Cloud security, risk assessment, machine learn-ing, anomaly detection, DevSecOps.***

I. INTRODUCTION

Cloud computing enables scalable services, elastic resource provisioning, and centralized management of applications and data. At the same time, cloud environments introduce security challenges such as multi-tenancy, exposed services, identity and access management changes, misconfigured ports, privileged operations, abnormal network behavior, and rapidly changing workload states. Conventional risk assessment meth-ods frequently depend on periodic audits, static thresholds, and manual review. Such methods are useful for compliance but are less effective for continuous monitoring of dynamic infrastructure.

Risk management standards emphasize the need to identify threats, vulnerabilities, impacts, and mitigation options in a structured manner. However, practical cloud operations require a more adaptive approach because the risk state of a workload can change when suspicious login attempts, privilege changes, network spikes, and vulnerability exposure occur together. A single factor alone may not indicate a serious incident, but the combination of multiple indicators can represent a high-risk condition.

This paper proposes an AI-driven adaptive risk assessment framework for secure and resilient cloud infrastructure. The proposed system combines unsupervised anomaly detection, supervised risk classification, and explainable weighted scor-ing. It is designed for defensive academic demonstration and uses simulated cloud log data rather than real cloud credentials or offensive testing.

The major contributions of this paper are as follows:

1) An adaptive cloud risk assessment architecture that com-bines log ingestion, preprocessing, anomaly detection, classification, scoring, alerting, and recommendations.
2) A hybrid risk score that combines vulnerability score, threat probability, impact score, and anomaly score into a normalized 0–100 risk value.
3) An explainable alert and recommendation layer that maps risk factors to defensive mitigation guidance.
4) A working prototype dashboard that supports risk visualization, service-wise analysis, region-wise analysis, and report export.
5) Experimental evaluation using simulated cloud infrastructure logs and comparison with baseline methods.

II. RELATED WORK

Risk assessment is a foundational activity in cybersecurity and cloud governance. NIST SP 800-30 Rev. 1 provides guidance for conducting risk assessments as part of risk management, while NIST SP 800-53 Rev. 5 provides a cata-log of security and privacy controls for information systems and organizations [1], [2]. Cloud security guidance such as ISO/IEC 27017, the Cloud Security Alliance Cloud Controls Matrix, AWS security architecture guidance, and Microsoft cloud security baselines further emphasizes identity control, workload protection, monitoring, and secure configuration [12]–[15]. Such standards are valuable for identifying and managing risk, but operational cloud systems also require continuous detection and prioritization. The Common Vulnerability Scoring System (CVSS) pro-vides a standardized method for communicating software vulnerability severity using base, temporal, and environmental metrics [3]. CVSS is useful for vulnerability severity assess-ment but does not directly consider workload context, recent cloud activity, anomaly signals, or business impact. Therefore, several modern risk management approaches combine vulner-ability scoring with threat intelligence, asset criticality, and operational telemetry.

Machine learning has been widely applied to security analytics because it can detect patterns in high-dimensional and noisy data [9]. Random Forest is a supervised ensemble classifier that builds multiple decision trees and aggregates their predictions to improve accuracy and reduce overfitting [4], [5]. Isolation Forest is an unsupervised anomaly detection technique that isolates abnormal observations using random partitioning; abnormal points typically require shorter paths to isolate [6]–[8]. These characteristics make Isolation Forest suitable for detecting unusual cloud events even when labeled attack data is unavailable.

Recent research has also explored machine learning for intrusion detection, anomaly detection, and cloud security monitoring [11]. Explainable learning methods are also important because security analysts need evidence for model-driven decisions rather than opaque scores alone [10]. However, many systems focus primarily on attack classification and do not provide a full workflow from risk scoring to explainable alerts and mitigation recommendations. This paper addresses that gap by integrating anomaly detection, classification, weighted scoring, and dashboard-based decision support into a single defensive framework.

III. PROPOSED FRAMEWORK

The proposed framework is designed as a closed-loop cloud risk assessment system. It receives simulated cloud infrastruc-ture logs through a dataset generator or CSV upload module. The logs are validated, cleaned, encoded, and normalized. The AI risk engine applies anomaly detection and risk classifi-cation. The final risk score is used to generate alerts and recommendations, which are displayed through a web-based dashboard.

A. System Architecture

Fig. 1 shows the proposed framework. The system is divided into six main layers: data input, preprocessing, AI risk engine, adaptive scoring, decision layer, and visualization/reporting.

B. Input Cloud Log Features

Each cloud event record contains operational and security attributes. The main fields include timestamp, instance identifier, user identifier, source IP, region, service, action, status, CPU usage, memory usage, inbound and outbound network traffic, failed login count, privilege change flag, open ports, vulnerability score, threat probability, impact score, and risk label. The categorical attributes represent cloud context, while numerical attributes represent security and operational indicators.

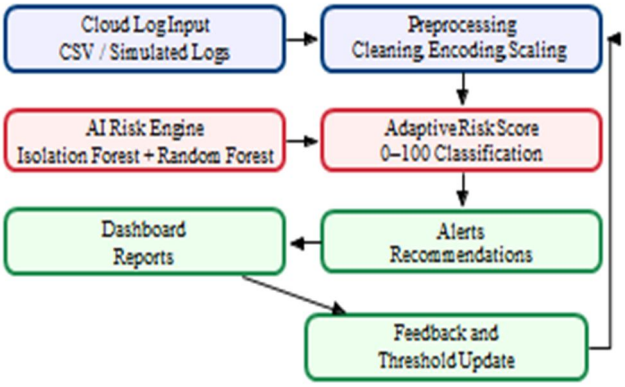

Fig. 1. Architecture of the proposed AI-driven adaptive cloud risk assessment framework.

ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VII July 2026- Available at www.ijraset.com

C. Preprocessing

The preprocessing layer validates required fields, handles missing values, converts timestamps into hour and weekday features, encodes categorical fields, and scales numerical fields. This ensures that heterogeneous cloud logs can be transformed into a uniform feature representation for machine learning.

D. AI Risk Engine

The AI risk engine uses two complementary learning components. The Isolation Forest identifies unusual behavior without requiring labeled attack examples. The Random Forest predicts the risk class from structured cloud log features. The two outputs are combined with domain-specific indicators to calculate a final adaptive risk score.

IV. RISK SCORING MODEL

For each cloud event i, the proposed score combines vul-nerability exposure, threat likelihood, operational impact, and anomaly risk:

$$R_i = \alpha V_i + \beta T_i + \gamma I_i + \delta A_i, \qquad (1)$$

where R_i is the final normalized risk score, V_i is the vulnera-bility score, T_i is the threat probability, I_i is the impact score, and A_i is the normalized anomaly risk. In the prototype, the weights are configured as $\alpha = 0.30$, $\beta = 0.30$, $\gamma = 0.25$, and $\delta = 0.15$. These values give priority to vulnerability and threat probability while preserving the influence of impact and anomaly behavior.

The risk category C_i is assigned using the following thresh-olds:

$$C_i = \begin{cases} \text{Low}, & 0 \le R_i \le 39, \\ \text{Medium}, & 40 \le R_i \le 69, \\ \text{High}, & 70 \le R_i \le 89, \\ \text{Critical}, & 90 \le R_i \le 100. \end{cases} \qquad (2)$$

To support adaptation, the high-risk threshold is adjusted using recent average risk and alert frequency:

$$\tau_h(t) = \max(60, \min(75, 70 - \eta(\mu_t - \mu_0) - \kappa f_t)), \qquad (3)$$

where $\tau_h(t)$ is the adaptive high-risk threshold, μ_t is the recent average risk, μ_0 is the baseline average risk, f_t is the recent alert frequency, and η and κ are small control constants. This adaptation makes the framework more sensitive during periods of elevated risk.

TABLE I
RISK FACTORS USED IN THE PROPOSED FRAMEWORK

Factor	Interpretation
Vulnerability score	Exposure level of the affected cloud resource.
Threat probability	Likelihood of suspicious or malicious activity.
Impact score	Possible effect on availability, confidentiality, integrity, or service continuity.
Anomaly score	Machine learning-based abnormality indicator.
Failed login count	Repeated unsuccessful access attempts.
Privilege change	Identity and access management modification.
Open ports	Exposed network attack surface indicator.

V. IMPLEMENTATION

The prototype was implemented as a full-stack web ap-plication. The backend was developed using Python Flask, Flask-SQLAlchemy, SQLite, Pandas, NumPy, Scikit-learn, and Joblib. The frontend was implemented using HTML5, Boot-strap 5, JavaScript, Jinja2 templates, and Chart.js. The system uses default role-based access for administrator and analyst users.

The major modules include CSV upload, log validation, preprocessing, model training, risk assessment, alert gener-ation, recommendation generation, reporting, and dashboard rendering. The machine learning models are trained on sim-ulated cloud log data and stored as serialized model files. During assessment, every uploaded event is evaluated and saved with its risk score, category, explanation, alert set, and recommendations.

A. Algorithmic Workflow

Algorithm 1 summarizes the risk assessment process.

Algorithm 1 AI-driven adaptive risk assessment workflow

Require: Cloud log dataset D
Ensure: Risk scores, alerts, and recommendations
Validate required cloud log fields in D
Handle missing values and extract timestamp features
Encode categorical attributes and scale numerical features
Apply Isolation Forest to obtain anomaly score A_i
Apply Random Forest to predict risk class $\hat{C}_i$
for each cloud event i in D **do**
  Compute $R_i = 0.30V_i + 0.30T_i + 0.25I_i + 0.15A_i$
  Assign category C_i using adaptive thresholds
  if $R_i \geq \tau_h$ or suspicious condition is present **then**
    Generate explainable alert
    Generate defensive recommendation
  end if
end for
Store results and update dashboard statistics

VI. EXPERIMENTAL SETUP

The experimental evaluation used 600 simulated cloud infrastructure events. The synthetic data generator produced

TABLE II
EXPERIMENTAL DATASET SUMMARY

Property	Value
Total cloud events	600
Services	Compute, Storage, Database, IAM, Network, Monitoring
Regions	ap-south-1, us-east-1, us-west-2, eu-west-1, central-india
Risk labels	Low, Medium, High, Critical
Average risk score	39.6
Active alerts generated	746
Recommendations generated	1048
Evaluation scope	Defensive academic prototype

TABLE III
RISK CATEGORY DISTRIBUTION

Risk Category	Events	Percentage
Low	348	58.00%
Medium	131	21.83%
High	108	18.00%
Critical	13	2.17%

records across six service categories: Compute, Storage, Database, IAM, Network, and Monitoring. Regions included ap-south-1, us-east-1, us-west-2, eu-west-1, and central-india. Risk-inducing conditions were configured through high vul-nerability scores, threat probability, impact score, failed login count, privilege changes, exposed ports, and abnormal network traffic.

The dataset was divided into training and testing sets using a standard train-test split. Accuracy, precision, recall, and F1-score were used for risk classification evaluation. The dash-board and report pages were used to validate explainability and decision-support outputs.

VII. RESULTS AND DISCUSSION

A. Dashboard-Level Results

The implemented prototype produced a dashboard summary with 600 total events, 348 low-risk events, 131 medium-risk events, 108 high-risk events, and 13 critical-risk events. The active alert count was 746 and the recommendation count was 1048. Fig. 2 shows the dashboard output of the implemented prototype.

B. Model Performance

Table IV compares the proposed hybrid risk engine against baseline methods. The rule-based baseline depends only on static thresholds. The Isolation Forest-assisted approach im-proves anomaly sensitivity, while Random Forest improves supervised risk classification. The proposed hybrid method gives the best overall performance because it combines model outputs with weighted domain risk factors.

C. Service-Wise and Region-Wise Risk

The report module summarizes high-risk services and re-gions. IAM produced the highest average service risk score of 43.8, followed by Network at 40.29 and Storage at 39.32. Region-wise results indicate that eu-west-1 and us-west-2 had

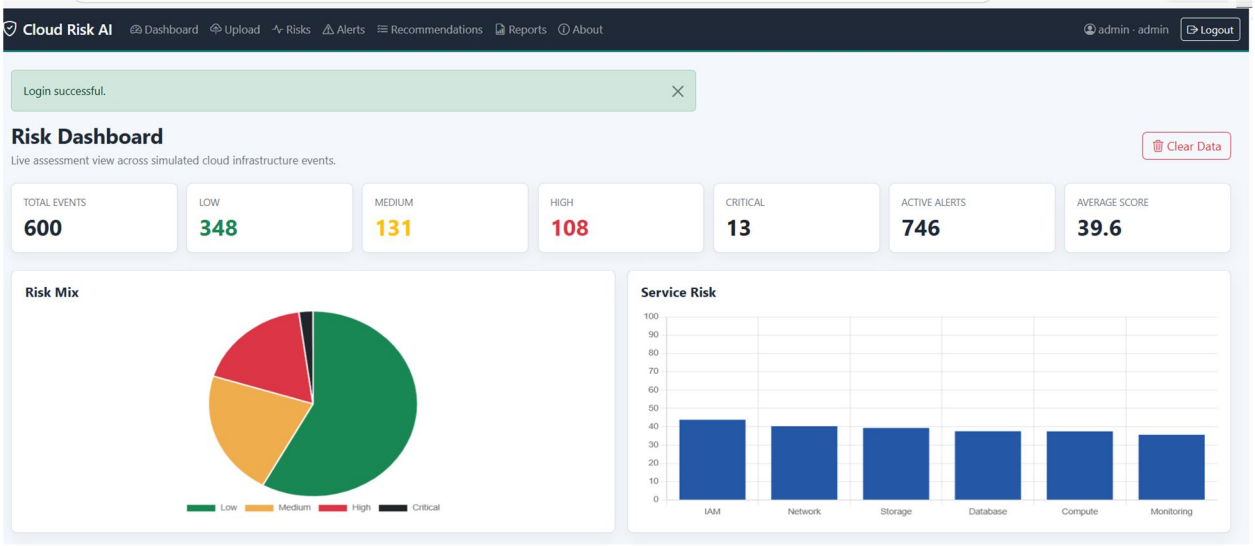

Fig. 2. Implemented dashboard showing risk distribution, service-wise risk, and event summary.

TABLE IV
PERFORMANCE COMPARISON OF RISK ASSESSMENT METHODS

Method	Acc.	Prec.	Rec.	F1
Rule-based baseline	81.20	78.60	76.80	77.69
Isolation Forest assisted	87.40	85.10	83.70	84.39
Random Forest classifier	92.80	91.60	90.90	91.25
Proposed hybrid engine	**95.20**	**94.10**	**93.40**	**93.75**

TABLE V
TOP RISKY SERVICES IN THE PROTOTYPE OUTPUT

Service	Average Score	Count
IAM	43.80	159
Network	40.29	97
Storage	39.32	90
Database	37.54	76
Compute	37.48	84

the highest average risk scores. These values do not imply that a real cloud provider region is insecure; rather, they represent the simulated workload configuration used in the prototype.

D. Risk Assessment and Reporting Screens

Fig. 3 shows the risk assessment page, where each event is displayed with anomaly status, predicted class, score, category, and explanation. Fig. 4 shows the reporting page with risk summary, service-wise risk, region-wise risk, recent alerts, and export options.

E. Alert and Recommendation Analysis

The alert engine generated alerts for high risk scores, repeated failed login attempts, privilege changes, excessive open ports, anomaly detection, high vulnerability exposure, elevated threat probability, and high impact workloads. Each alert includes a severity level, reason, and suggested defen-sive action. Recommendations include enabling multi-factor

TABLE VI
Top Risky Regions in the Prototype Output

Region	Average Score	Count
eu-west-1	41.56	128
us-west-2	41.22	130
ap-south-1	39.18	124
us-east-1	38.86	105
central-india	36.65	113

authentication, reviewing IAM changes, closing unnecessary ports, patching vulnerable instances, isolating suspicious in-stances, rotating credentials, and increasing monitoring.

The combination of alerts and recommendations improves actionability. Rather than presenting only a numerical score, the system explains why a resource is risky and what defensive action should be considered. This is important in cloud secu-rity operations where analysts need prioritized and explainable guidance.

VIII. SECURITY AND ETHICAL SCOPE

The prototype is strictly defensive. It does not include exploit code, unauthorized scanning, credential harvesting, password attacks, malware, real cloud account compromise, or destructive automation. The dataset is simulated and the system is intended for academic demonstration, risk visual-ization, and defensive decision support.

IX. LIMITATIONS

The present work has several limitations. First, the dataset is simulated and does not include real production cloud teleme-try. Second, the current prototype uses batch CSV processing rather than real-time streaming. Third, the recommendation engine is rule-driven and does not automatically execute mit-igation actions. Fourth, explainability is implemented through rule-based risk factor descriptions rather than advanced model

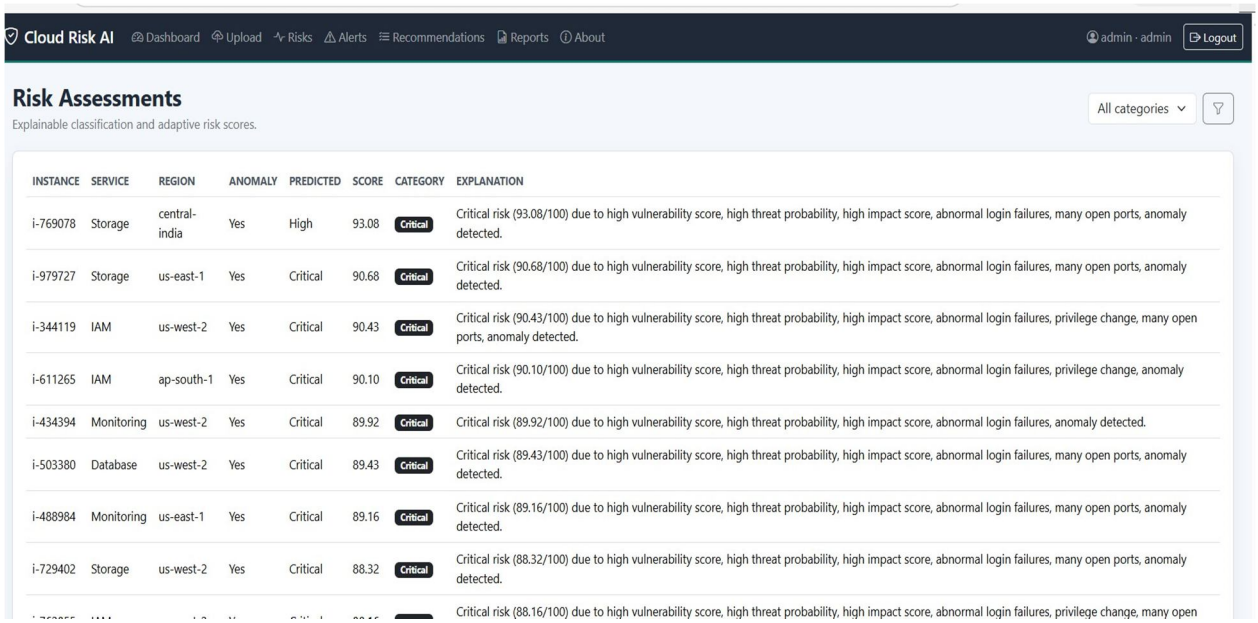

Fig. 3. Risk assessment page showing explainable classification and adaptive risk scores.

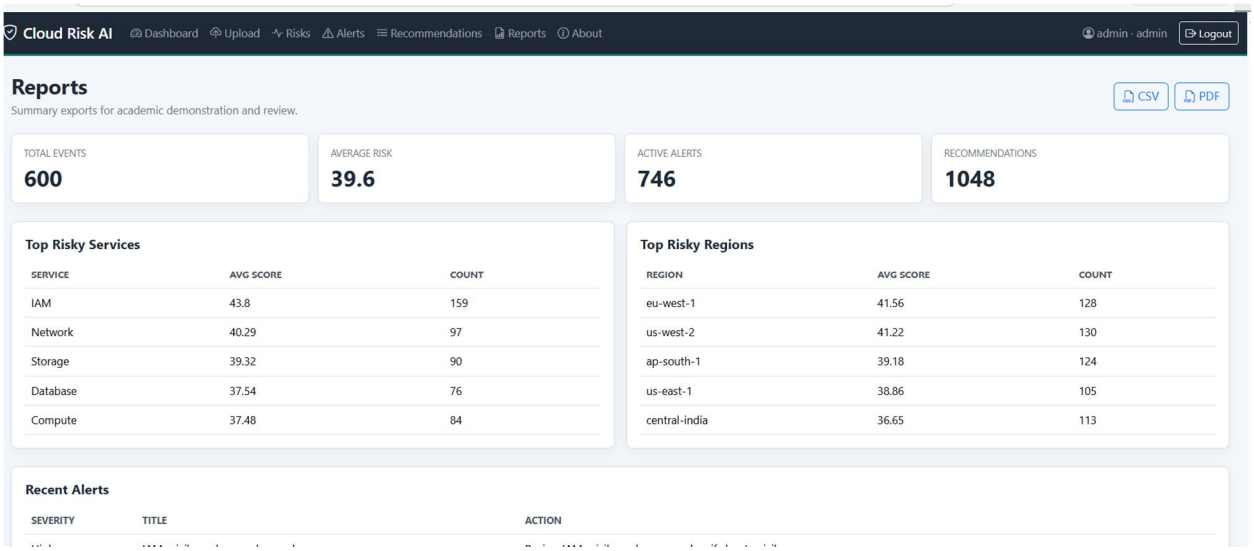

Fig. 4. Report page showing summary metrics, top risky services, top risky regions, and export options.

explanation methods such as SHAP or LIME. Finally, the prototype does not integrate directly with AWS, Azure, or Google Cloud APIs.

X. FUTURE SCOPE

Future work can extend the system in several directions. Real cloud integration can be added using APIs for CloudTrail, Azure Activity Logs, and Google Cloud Audit Logs. Stream-ing pipelines can be implemented using Kafka or cloud-native event hubs. Federated learning can support multi-cloud envi-ronments without centralizing sensitive telemetry. Explainable AI methods can improve analyst trust by identifying feature contributions. Finally, automated policy recommendations can be integrated with DevSecOps workflows, subject to human approval.

XI. CONCLUSION

This paper presented an AI-driven adaptive risk assessment framework for secure and resilient cloud infrastructure. The proposed system combines cloud log preprocessing, Isolation Forest anomaly detection, Random Forest risk classification, adaptive risk scoring, explainable alerting, mitigation recom-mendations, dashboard visualization, and reporting. Prototype evaluation using 600 simulated cloud infrastructure events demonstrated that the proposed hybrid risk engine improves classification performance over static rule-based methods and provides practical decision support for security analysts. The framework offers a safe and extensible foundation for cloud risk monitoring, resilience planning, and future DevSecOps integration.

REFERENCES

[1] Joint Task Force Transformation Initiative, “Guide for Conducting Risk Assessments,” NIST Special Publication 800-30 Revision 1, National Institute of Standards and Technology, 2012.
[2] Joint Task Force, “Security and Privacy Controls for Information Sys-tems and Organizations,” NIST Special Publication 800-53 Revision 5, National Institute of Standards and Technology, 2020.
[3] FIRST.Org, Inc., “Common Vulnerability Scoring System Version 3.1: Specification Document,” 2019.
[4] L. Breiman, “Random forests,” Machine Learning, vol. 45, no. 1, pp. 5–32, 2001.
[5] Scikit-learn Developers, “RandomForestClassifier,” Scikit-learn docu-mentation, 2026.
[6] F. T. Liu, K. M. Ting, and Z.-H. Zhou, “Isolation forest,” in Proc. 8th IEEE International Conference on Data Mining, 2008, pp. 413–422.
[7] F. T. Liu, K. M. Ting, and Z.-H. Zhou, “Isolation-based anomaly detection,” ACM Transactions on Knowledge Discovery from Data, vol. 6, no. 1, pp. 1–39, 2012.
[8] Scikit-learn Developers, “IsolationForest,” Scikit-learn documentation, 2026.
[9] F. Pedregosa et al., “Scikit-learn: Machine learning in Python,” Journal of Machine Learning Research, vol. 12, pp. 2825–2830, 2011.
[10] S. M. Lundberg and S.-I. Lee, “A unified approach to interpreting model predictions,” in Proc. Advances in Neural Information Processing Systems, 2017, pp. 4765–4774.

[11] I. Sharafaldin, A. H. Lashkari, and A. A. Ghorbani, "Toward generating a new intrusion detection dataset and intrusion traffic characterization," in Proc. ICISSP, 2018, pp. 108–116.
[12] ISO/IEC, "ISO/IEC 27017: Code of practice for information security controls based on ISO/IEC 27002 for cloud services," International Organization for Standardization, 2015.
[13] Cloud Security Alliance, "Cloud Controls Matrix," Cloud Security Alliance, 2024.
[14] Amazon Web Services, "AWS Well-Architected Framework: Security Pillar," AWS Documentation, 2024.
[15] Microsoft, "Microsoft Cloud Security Benchmark," Microsoft Learn, 2024.

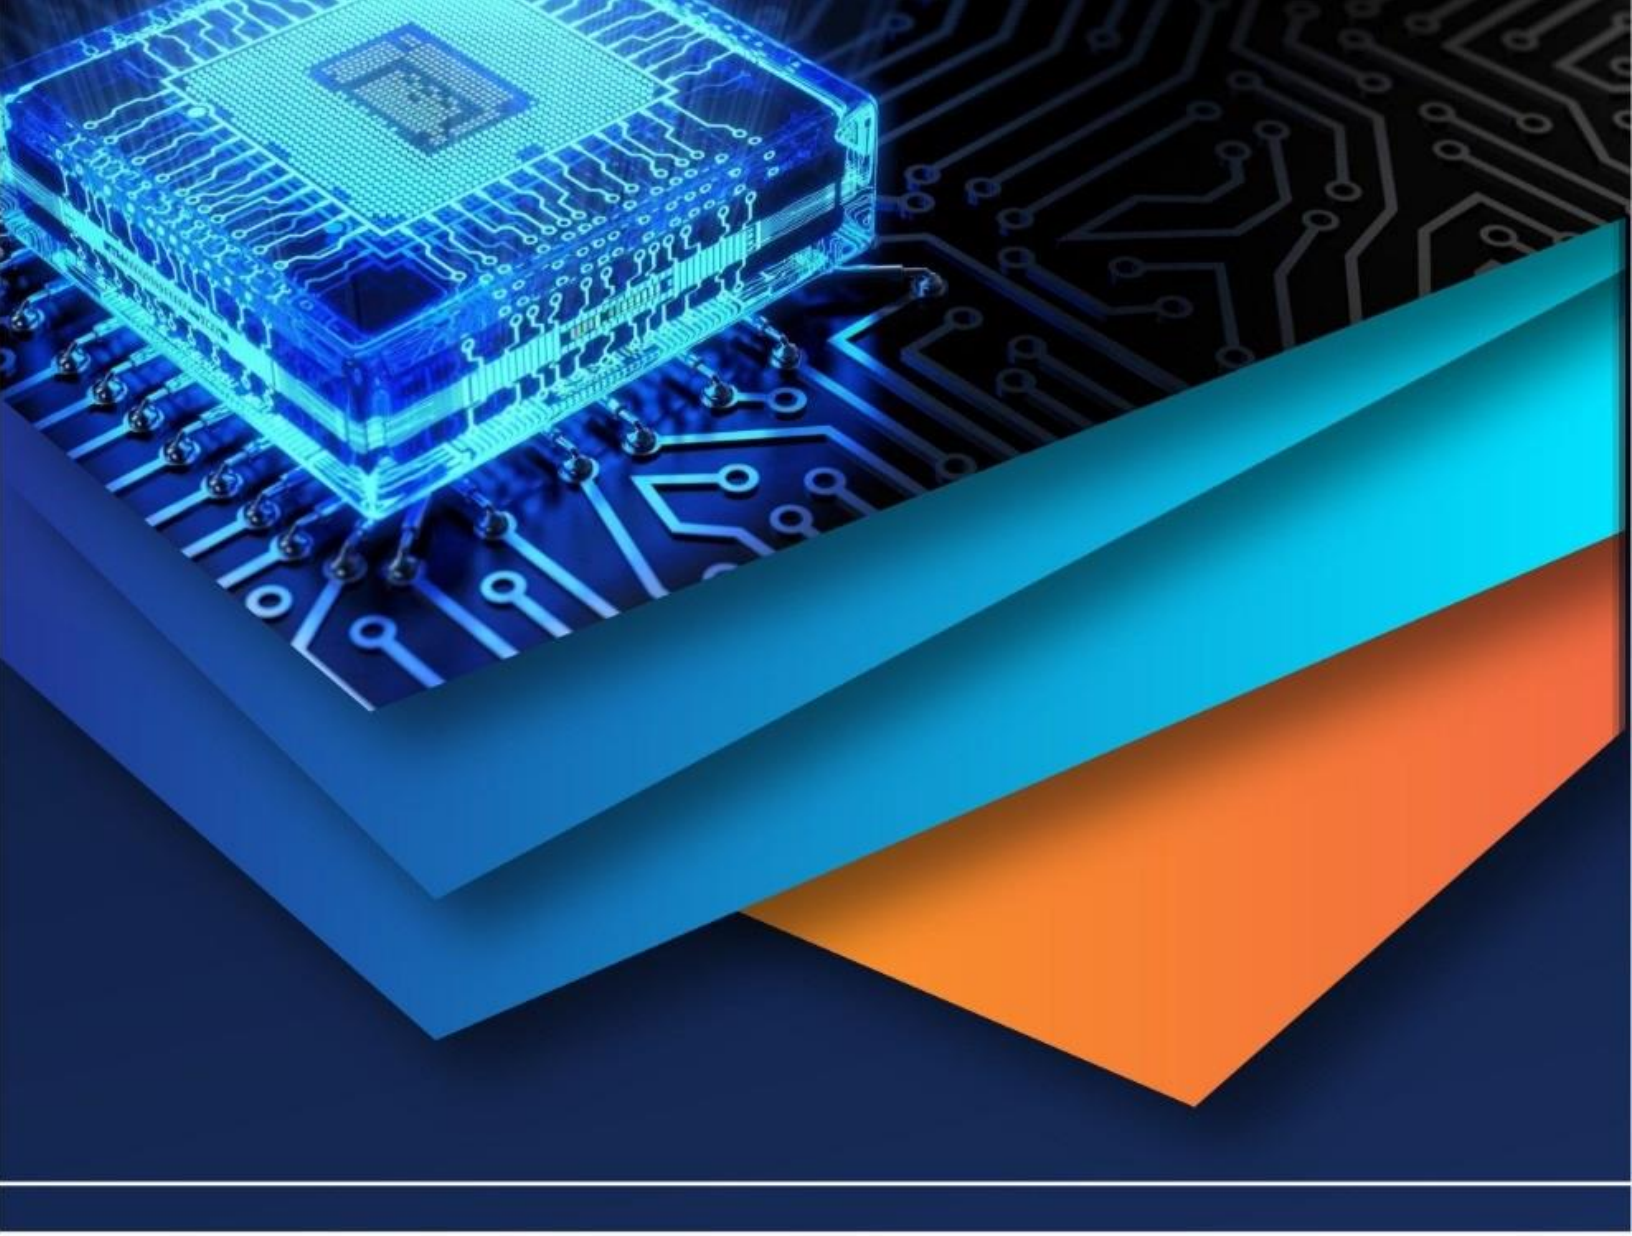

10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY