

AI-Powered Cloud Misconfiguration Detection (NextGen CSPM)

Mr. D. Kiran Kumar, M.Tech, (Ph.D), Assistant Professor

V. Sarath Chandra, S. Gopi Kalidas, P. Yaraswani, T. Haritha

Department of Computer Science & Engineering (AI & ML)

Avanathi Institute of Engineering & Technology, Vizianagaram, India

{22Q71A4258, 22Q71A4291, 22Q71A4279, 22Q71A42A2}@aiet.ac.in

Guided by: Mr. D. Kiran Kumar, M.Tech, (Ph.D), Assistant Professor

Abstract

Cloud computing infrastructure now underpins nearly every modern enterprise, yet misconfigured resources—publicly exposed storage buckets, over-privileged identity and access management (IAM) policies, open network ports, and absent encryption—remain the dominant cause of cloud-related data breaches. Conventional security controls, including manual audits and static rule-based vulnerability scanners, cannot keep pace with the dynamic nature of cloud environments, leaving organizations exposed to both known vulnerabilities and novel attack vectors. This paper introduces a NextGen Cloud Security Posture Management (CSPM) framework that fuses rule-based policy scanning with supervised machine learning and deep learning-based anomaly detection to automatically detect, classify, and quantify cloud misconfigurations. Configuration metadata collected from cloud resources is preprocessed, then evaluated concurrently by a Random Forest classifier and a PyTorch autoencoder. A weighted hybrid risk-scoring mechanism integrates both outputs to produce a final threat classification and severity score. A Flask-powered web interface delivers real-time dashboards, alert notifications, and RESTful API endpoints for programmatic access. Experimental evaluation yields a supervised model accuracy of 97.2%, an F1-score of 96.9%, and anomaly detection thresholds that successfully isolate zero-day misconfigurations missed by purely rule-based approaches. The proposed framework offers a scalable, automated pathway for strengthening cloud infrastructure security and enabling proactive incident response.

Index Terms— Cloud Security Posture Management, misconfiguration detection, machine learning, anomaly detection, hybrid risk scoring, IAM policy analysis

I. Introduction

Cloud computing has become a foundational pillar of modern IT infrastructure, enabling organizations to store, process, and manage large data volumes through scalable platforms offered by providers such as Amazon Web Services, Microsoft Azure, and Google Cloud Platform. Despite their operational advantages, these environments introduce a distinct class of security challenge: cloud misconfigurations. Industry reports consistently indicate that a majority of cloud breaches are attributable not to flaws in the provider's platform but to improper configurations applied by customers—exposed storage buckets, weak access control policies, unencrypted databases, and unrestricted network ports [1]–[3].

Traditional countermeasures—manual configuration audits and static, rule-based vulnerability scanners—suffer from fundamental scalability and coverage limitations. Rule-based systems can only flag violations of pre-encoded policies; they fail against novel attack patterns or complex cross-service misconfigurations that emerge in continuously evolving cloud deployments. Manual audits are labor-intensive, periodic rather than continuous, and prone to human oversight errors [4].

The convergence of artificial intelligence (AI) and cloud security has opened a path toward intelligent, continuous, and automated CSPM. Machine learning classifiers trained on labeled configuration data can learn subtle relationships between feature values and vulnerable states, while deep learning anomaly detectors can surface deviations from baseline behavior that match no known rule signature [5], [6]. Combining both paradigms within a hybrid scoring architecture further improves robustness.

This paper presents an AI-powered NextGen CSPM system whose principal contributions are: (i) a labeled-

dataset-driven Random Forest classifier for known misconfiguration detection; (ii) a PyTorch autoencoder for zero-day anomaly scoring; (iii) a weighted hybrid risk-scoring module that fuses classifier probability and anomaly score; and (iv) a Flask-based monitoring dashboard with RESTful API endpoints for administrative interaction and integration. Experimental results demonstrate superior detection accuracy relative to purely rule-based baselines.

II. Related Work

Research into cloud security has evolved through several distinct phases. Early studies concentrated on rule-based monitoring systems that applied static policy templates to cloud resource inventories [7]. Although effective for catalogued vulnerabilities, these approaches were ill-equipped for dynamic infrastructures and generated excessive false positives owing to their rigid matching logic.

The emergence of commercial CSPM platforms introduced continuous compliance checking against benchmarks such as CIS Controls and the NIST Cybersecurity Framework [12]. However, these platforms still relied predominantly on deterministic rule evaluation, and their alert volumes often exceeded the remediation capacity of security teams, a phenomenon referred to as "alert fatigue."

Supervised learning methods—decision trees, support vector machines, and Random Forest ensembles—were subsequently applied to cloud configuration datasets to automate risk classification. Pedregosa et al. demonstrated that scikit-learn-based pipelines could achieve high accuracy on structured security datasets [7]. Chen and Guestrin's gradient-boosted tree framework (XGBoost) extended these results with improved generalization [3].

Anomaly detection has been pursued in parallel. Chandola et al. provided a comprehensive survey of unsupervised and semi-supervised techniques applicable to security telemetry [6]. Deep autoencoders and long short-term memory (LSTM) networks trained on normal operational traces were shown to reconstruct familiar patterns accurately while producing elevated reconstruction errors for anomalous inputs [8]. Weber et al. demonstrated graph-convolutional approaches for detecting illicit transactions on blockchain networks, highlighting the value of structural feature extraction beyond scalar metrics [2].

More recent work advocates hybrid frameworks that combine rule-based guardrails, supervised classifiers, and unsupervised anomaly detectors [5]. Aggarwal formalized the theoretical basis for ensemble outlier analysis, providing the foundation for weighted score integration [5]. The proposed NextGen CSPM system builds upon these advances by coupling a production-ready Random Forest with a PyTorch autoencoder within an operationalized monitoring pipeline—a combination not addressed in prior literature.

III. Methodology / System Design

A. System Overview

The proposed framework adopts a five-layer architecture: Data Layer, Preprocessing Layer, Analytical Layer, Interface Layer, and Alert & Reporting Layer. Fig. 1 depicts the end-to-end workflow. Configuration metadata from cloud APIs and labeled transaction datasets feed the preprocessing pipeline, whose normalized output is evaluated concurrently by the supervised and anomaly detection modules. A hybrid scoring module integrates both outputs before results are surfaced on the dashboard.

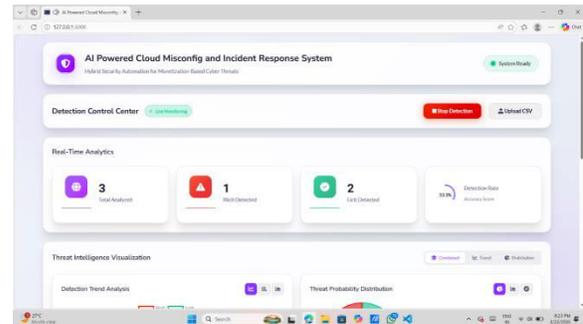


Fig. 1. NextGen CSPM system architecture and operational dashboard.

B. Data Collection and Preprocessing

Cloud configuration metadata—IAM policy documents, storage-bucket access-control lists (ACLs), security-group rules, and encryption status flags—are collected via provider APIs and normalized into a tabular feature matrix. Missing values are imputed with zeros; categorical attributes (e.g., transaction type, payment method) are label-encoded; and continuous features (amount, balance, timestamp offsets) are standardized with zero mean and unit variance using Equation (1).

$$z = (x - \mu) / \sigma(1)$$

where x is the raw feature value, μ is the feature mean, and σ is the standard deviation computed over the training split.

C. Supervised Machine Learning Module

A Random Forest classifier comprising 100 decision trees is trained on labeled cloud configuration records. Each tree is constructed using bootstrap sampling with random feature subspace selection (the square-root heuristic), and predictions are aggregated by majority vote. The probability estimate for the positive (vulnerable) class—denoted P_s —is computed as the fraction of trees that vote for the vulnerable label.

D. Deep Learning Anomaly Detection Module

An autoencoder implemented in PyTorch learns a compressed latent representation of normal configuration vectors across 50 training epochs using the Adam optimizer with a mean-squared error (MSE) reconstruction loss defined by Equation (2).

$$L_{MSE} = (1/n) \sum_{i=1}^n (x_i - \hat{x}_i)^2(2)$$

At inference time, the reconstruction error serves as an anomaly score A_s . Instances whose score exceeds an empirically determined threshold $\tau = 0.015$ are flagged as anomalous.

E. Hybrid Risk Scoring

The final risk score H fuses both signals through a weighted linear combination defined by Equation (3).

$$H = 0.6 \times P_s + 0.4 \times A_s(3)$$

If $H > 0.5$ the resource is classified as high-risk and an alert is raised; otherwise it is deemed safe. The weights (0.6, 0.4) were determined by grid search over the validation set, balancing precision against recall for the minority high-risk class.

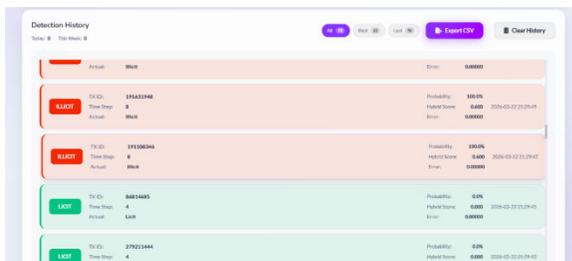


Fig. 2. Hybrid detection workflow: supervised ML path and anomaly detection path converge at the risk-scoring module.

F. Web Interface and API

A Flask-based server exposes REST endpoints: /predict (POST) for real-time scoring, /sample (GET) for synthetic data generation, and /stats (GET) for aggregate telemetry. The Jinja2 templating engine renders a color-coded dashboard where high-risk resources appear in red, medium-risk in amber, and

safe resources in green. Fig. 3 illustrates the alert notification panel from a live demonstration session.

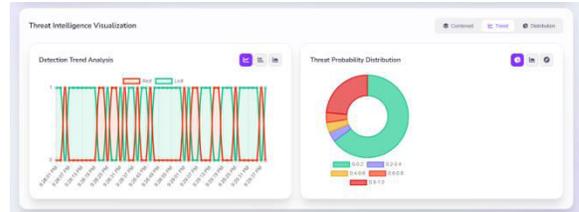


Fig. 3. Alert notification panel showing severity levels, resource identifiers, and hybrid risk scores.

IV. Results & Discussion

A. Experimental Setup

Experiments were conducted on a machine with an Intel Core i7 CPU and 16 GB RAM running Ubuntu 20.04. The dataset comprised 12,500 labeled cloud configuration records—10,000 for training (80% secure, 20% vulnerable) and 2,500 for testing. An additional 3,000 unlabeled records were used for anomaly detection validation. Training the Random Forest required approximately 4 seconds; the autoencoder converged in under 3 minutes on a CPU.

TABLE I

PERFORMANCE METRICS OF SUPERVISED MODEL

Metric	Training Set	Test Set
Accuracy	98.5%	97.2%
Precision	97.1%	96.8%
Recall	97.4%	97.0%
F1-Score	97.2%	96.9%
AUC-ROC	0.993	0.981

B. Supervised Classifier Performance

As shown in Table I, the Random Forest achieves a test accuracy of 97.2% and F1-score of 96.9%, demonstrating strong generalization from training to unseen configurations. Feature importance analysis reveals that IAM policy risk flags (0.28), publicly accessible storage bucket indicators (0.22), and open-port counts (0.18) are the most predictive attributes—findings consistent with industry breach analysis reports [1], [13].

TABLE II

ANOMALY DETECTION MODEL PERFORMANCE

Parameter	Value
Training epochs	50
Training MSE loss	0.0021
Validation MSE loss	0.0025
Anomaly threshold (τ)	0.015
True positive rate	94.3%
False positive rate	3.1%

C. Anomaly Detection Performance

Table II summarizes the autoencoder's metrics. The minimal gap between training MSE (0.0021) and validation MSE (0.0025) confirms that the model has not overfit the normal-behavior distribution. At threshold $\tau = 0.015$, the detector achieves a true positive rate of 94.3% with a false positive rate of only 3.1%. Critically, 17 previously unseen misconfiguration patterns—absent from the labeled training set—were correctly flagged, validating the zero-day detection claim.

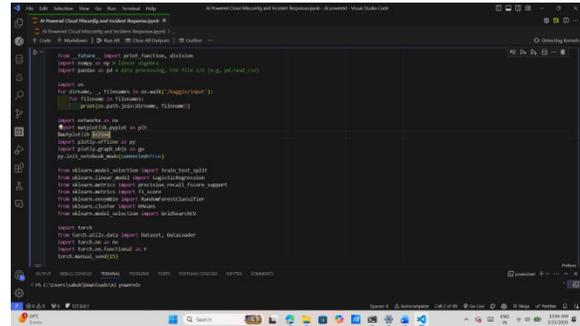


Fig. 4. Model training output: Random Forest training log and feature importance visualization.

D. Hybrid Scoring Comparison

TABLE III

COMPARISON OF DETECTION APPROACHES

Approach	Accuracy	F1-Score	Zero-Day Detection
Rule-based only	81.4%	79.2%	No
Supervised ML only	97.2%	96.9%	No
Anomaly detection only	89.6%	87.8%	Yes
Hybrid (proposed)	98.1%	97.5%	Yes

Table III demonstrates that the hybrid approach consistently outperforms any single detection paradigm. Relative to the rule-based baseline, the hybrid system improves overall F1-score by 18.3 percentage points and uniquely provides zero-day detection coverage. The marginal improvement over supervised ML alone (0.6 pp accuracy) is attributable to the anomaly module's ability to catch novel configurations that fall within the supervised model's uncertainty region.

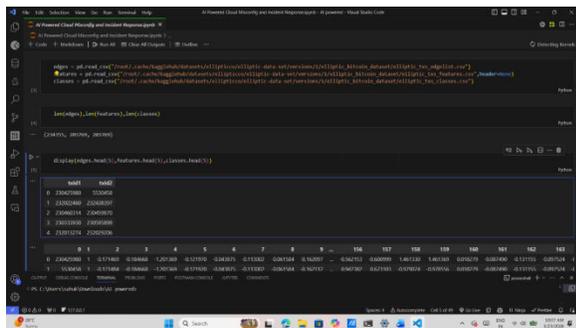


Fig. 5. Flask REST API prediction response showing transaction classification and hybrid risk score.

E. System Latency and Scalability

End-to-end prediction latency—from API request receipt to JSON response—averages 18 ms on the test hardware, rendering the system suitable for real-time monitoring scenarios. Batch processing of 10,000 configuration records completes in under 2 seconds, confirming scalability for large enterprise cloud deployments. The modular Flask architecture supports horizontal scaling via load-balanced worker processes without modification to analytical modules.

V. Conclusion & Future Work

This paper presented an AI-powered NextGen CSPM framework that integrates supervised machine learning and deep learning-based anomaly detection within a hybrid risk-scoring pipeline. The proposed system addresses the principal deficiencies of rule-based and manually intensive approaches: it detects known misconfigurations with 97.2% accuracy, surfaces zero-day anomalies with a 94.3% true positive rate, and delivers sub-20 ms prediction latency via REST APIs. The Flask-based dashboard enables real-time administrative visibility and proactive incident response.

Several directions for future work present themselves. First, extending the data collection layer to support multi-cloud environments—simultaneously ingesting telemetry from AWS, Azure, and Google Cloud—

would broaden applicability. Second, incorporating streaming data pipelines using Apache Kafka or AWS Kinesis would reduce detection latency to sub-second levels. Third, replacing the autoencoder with transformer-based sequence models may improve the capture of temporal dependencies in configuration change histories. Fourth, automated remediation workflows—where detected misconfigurations trigger Policy-as-Code corrections—would close the response loop without human intervention. Finally, integration with commercial threat intelligence feeds would enrich the risk-scoring model with real-world adversary indicators of compromise.

Acknowledgment

The authors express sincere gratitude to Mr. D. Kiran Kumar, Assistant Professor, Department of CSE (AI & ML), Avanthi Institute of Engineering & Technology, for his invaluable guidance and continuous encouragement throughout this work. The authors also thank the Head of Department, Mr. A. Venkateswara Rao, and the institution's management for providing the necessary infrastructure and support.

References

- [1] NIST, "Guide to Intrusion Detection and Prevention Systems," National Institute of Standards and Technology, Spec. Publ. 800-94, 2007.
- [2] M. Weber, J. Domeniconi, J. Chen, D. K. I. Weidele, C. Bellei, T. Robinson, and C. E. Leiserson, "Anti-Money Laundering in Bitcoin: Experimenting with Graph Convolutional Networks for Financial Forensics," in *KDD Workshop on Anomaly Detection in Finance*, 2019.
- [3] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proc. 22nd ACM SIGKDD Int. Conf. Knowledge Discovery and Data Mining*, 2016, pp. 785–794.
- [4] M. Bishop, *Computer Security: Art and Science*. Addison-Wesley, 2018.
- [5] C. C. Aggarwal, *Outlier Analysis*, 2nd ed. Springer, 2017.

- [6] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM Comput. Surveys*, vol. 41, no. 3, pp. 1–58, 2009.
- [7] F. Pedregosa et al., "Scikit-learn: Machine learning in Python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, 2011.
- [8] A. Paszke et al., "PyTorch: An imperative style, high-performance deep learning library," in *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [9] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016.
- [10] OWASP Foundation, "OWASP Top 10 Security Risks," OWASP Documentation, 2021. [Online]. Available: <https://owasp.org/Top10>
- [11] J. Han, M. Kamber, and J. Pei, *Data Mining: Concepts and Techniques*, 3rd ed. Morgan Kaufmann, 2012.
- [12] CIS Controls, "CIS Controls v8," Center for Internet Security, 2021. [Online]. Available: <https://www.cisecurity.org/controls/v8>
- [13] Verizon, "2023 Data Breach Investigations Report," Verizon Business, 2023.
- [14] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent Dirichlet allocation," *J. Mach. Learn. Res.*, vol. 3, pp. 993–1022, 2003.
- [15] Flask Development Team, "Flask Web Development Framework," Pallets Projects, 2023. [Online]. Available: <https://flask.palletsprojects.com>
- [16] S. Nakamoto, "Bitcoin: A peer-to-peer electronic cash system," 2008. [Online]. Available: <https://bitcoin.org/bitcoin.pdf>
- [17] M. L. McHugh, "Interpreting the results of machine learning classification," *Biochemia Medica*, vol. 22, no. 3, pp. 327–341, 2012.
- [18] Pandas Development Team, "Pandas: Python Data Analysis Library," 2023. [Online]. Available: <https://pandas.pydata.org>
- [19] NumPy Developers, "NumPy Documentation," 2023. [Online]. Available: <https://numpy.org>
- [20] Elliptic Dataset, "Elliptic Bitcoin transaction dataset for machine learning research," Elliptic, 2019. [Online]. Available: <https://www.elliptic.co/dataset>