Published by :
http://www.ijert.org

International Journal of Engineering Research & Technology (IJERT)
ISSN: 2278-0181
Vol. 14 Issue 05, May-2025

AI-Augmented Threat Detection and Policy Drift Remediation in Hybrid Cloud Network Security Architectures

Kapil C. Wannere
Individual Researcher - USA

Hybrid Cloud, Threat Detection, Policy Drift, Artificial Intelligence, Machine Learning, Network Security, Infrastructure-as-Code, Azure, GCP, Automation, Zero Trust

I. INTRODUCTION

Abstract

The increasing complexity and scale of hybrid cloud environments—spanning private data centers, Microsoft Azure, and Google Cloud Platform (GCP)—has significantly strained traditional network security controls. Static firewalls, manually curated policies, and signature-based detection models are insufficient in identifying evolving threats and preventing misconfiguration-induced vulnerabilities. This paper presents an AI-augmented framework that integrates real-time telemetry analysis, machine learning-based threat detection, and automated policy drift remediation. Our solution leverages Infrastructure-as-Code (IaC) to continuously validate and enforce secure baselines while dynamically adapting to threats and configuration deviations. Experimental evaluation in a hybrid lab testbed demonstrates improved detection accuracy, accelerated response times, and reduced operational overhead. The proposed system offers a scalable, resilient, and intelligent foundation for securing hybrid cloud networks.

Keywords

Hybrid cloud computing offers organizations flexibility, scalability, and cost savings by combining on-premises infrastructure with public cloud services such as Microsoft Azure and Google Cloud Platform (GCP). However, this architectural model also creates fragmented visibility, inconsistent enforcement of policies, and heightened risk exposure due to dynamic and distributed environments. These complexities require automated detection systems capable of handling distributed telemetry, interpreting diverse control plane signals, and performing risk-aware policy adjustments. The hybrid cloud model demands that security policies be both platform-agnostic and context-aware, enabling consistent enforcement regardless of the underlying infrastructure provider. Our proposed system addresses these requirements through modular telemetry collection, AI-driven decision engines, and codified policy enforcement mechanisms.

Traditional perimeter security models assume static configurations, fixed IP addresses, and deterministic network paths—assumptions which break down in cloud-native architectures. In such setups, policy drift is a major concern. Policy drift occurs when the actual security configuration diverges from the intended policy baseline—either due to manual errors, automated deployment changes, or overlooked third-party integrations.

Furthermore, the threat landscape continues to evolve with attackers leveraging lateral movement, credential misuse, and misconfigured access control rules to bypass security defenses. Security teams face alert fatigue and increasing complexity in correlating multi-source telemetry to understand the root cause of incidents. Artificial intelligence (AI), particularly machine learning (ML), presents a powerful solution by enabling automated pattern recognition, anomaly detection, and intelligent remediation decisions across large-scale hybrid cloud networks.

II. RELATED WORK

Prior efforts in hybrid cloud security have largely focused on centralized monitoring through Security Information and Event Management (SIEM) systems such as Azure Sentinel, Splunk, or Chronicle. While useful for log aggregation and alerting, these platforms lack the ability to interpret contextual policy intent or automate corrective enforcement. Platforms like AWS Config and Azure Policy provide configuration compliance checks, but are rule-based and reactive.

Machine learning-based intrusion detection systems (IDS) like Zeek, Snort with AI extensions, and commercial NDR (Network Detection and Response) platforms have shown promise in detecting previously unseen attack patterns. Academic research has explored the use of autoencoders, isolation forests, and graph-based analytics for detecting anomalies in network traffic.

Our work extends beyond detection by integrating AI-based analysis with real-time remediation through Infrastructure-as-Code. Additionally, we introduce a drift validator component that proactively reconciles deviations using IaC baselines and integrates with SOAR (Security Orchestration, Automation and Response) for policy enforcement.

III. PROPOSED ARCHITECTURE

The proposed architecture includes four main layers: (1) Telemetry Ingestion, (2) AI Processing Engine, (3) Drift Validator, and (4) Remediation Orchestrator.

Telemetry is sourced from Azure NSG flow logs, GCP VPC logs, firewall rules, IAM changes, and routing updates. These logs are normalized and forwarded to the AI Engine, which detects anomalies using supervised and unsupervised models.

Simultaneously, a configuration snapshot is taken from deployed resources and compared against the intended policy baseline defined in Git-managed Terraform or Bicep templates. Any drift is flagged, scored for severity, and queued for remediation. The Remediation Orchestrator pushes IaC corrections back to the cloud environment and updates security operations dashboards.

The entire process is logged, version-controlled, and exposed via REST APIs and webhook integrations for audit and incident response.

IV. AI METHODOLOGY

The AI engine comprises a hybrid of anomaly detection and behavior-based classification models. The core includes the following modules:

- Anomaly Detection: Utilizes Isolation Forests and DBSCAN clustering to identify deviations in flow logs and access patterns.
- Time Series Prediction: Long Short-Term Memory (LSTM) networks forecast resource behavior (e.g., login frequency, traffic volume) and detect unexpected spikes or dips.
- Drift Scoring Model: Assesses severity based on delta between intended configuration and actual state using rule entropy, ACL divergence, and peer-group access anomalies.
- Reinforcement Learning Agent: Optimizes remediation decisions using a reward-based model trained on prior response outcomes.

The system is continuously retrained using synthetic datasets from simulated environments, as well as anonymized enterprise telemetry streams. Feature engineering includes packet entropy, port scan ratios, rule redundancy scores, and anomaly correlation graphs.

Additionally, we included a feature attribution module that utilizes SHAP (SHapley Additive exPlanations) to interpret predictions from the LSTM and Isolation Forest models. Hyperparameters were optimized using grid search and cross-validation over multiple telemetry slices. We also tested ensemble methods combining logistic regression and random forest classifiers for known policy drift signatures. The model evaluation pipeline was implemented using MLflow and integrated with a CI/CD system to support versioned model deployment.

V. IMPLEMENTATION AND EVALUATION

We implemented the framework in a lab environment comprising:

- Azure hub-and-spoke VNets
- GCP Shared VPCs
- Simulated on-premises network via EVE-NG with Cisco and Palo Alto appliances

Telemetry was gathered using Fluent Bit and forwarded to a Kafka stream. The AI engine was hosted on an AKS (Azure Kubernetes Service) cluster with GPU acceleration for LSTM training. Configuration states were versioned in a GitHub repo, and remediation executed through Terraform Cloud pipelines.

Tested threat scenarios:
1. NSG misconfiguration allowing open RDP (port 3389)
2. Abnormal east-west traffic spike across regions
3. Inactive accounts making privileged API calls
4. Drifted IAM roles exposing unused services

Evaluation Metrics:
- Threat Detection Precision: 94.1%
- Policy Drift Accuracy: 91.7%
- False Positive Rate: 4.6%
- Automated Remediation Success Rate: 89.3%
- Mean Time to Respond (MTTR): 42 seconds

The testbed included 150+ simulated nodes, including Windows and Linux VMs, containers running microservices, and API gateways across different cloud regions. Traffic generators simulated legitimate and malicious activity such as DDoS bursts, lateral movement across subnets, and slow exfiltration. Policy drift was introduced intentionally by modifying IaC baselines mid-deployment and comparing remediation latency against baseline SLAs. The system demonstrated consistent performance under high-load scenarios, processing over 1 million telemetry records daily with 98.7% inference uptime.

VI. RESULTS

The framework demonstrated robust performance across detection, drift identification, and automated enforcement tasks. During simultaneous attack simulation across cloud and on-premises resources, the AI system detected and contained 92% of threats within 30 seconds.
The drift validator correctly flagged 27 configuration anomalies out of 29 injected cases, highlighting high fidelity in IaC reconciliation.

The following table summarizes key results:

Metric	Score	Comments
Threat Detection Precision	94.1%	High accuracy with LSTM + Isolation Forest
Drift Detection Accuracy	91.7%	Detected subtle ACL misalignments
Remediation Success Rate	89.3%	Via Terraform Cloud pipelines
False Positive Rate	4.6%	Tuned threshold on policy scoring
Mean Time to Remediate	42s	Significant time reduction from baseline

VII. DISCUSSION

The key strength of the system lies in its integration of AI models with automated remediation logic tied to cloud-native tooling. Rather than relying solely on alert generation, the framework offers closed-loop policy correction, reducing human intervention and misconfiguration exposure windows. Challenges include model drift, particularly with rapidly evolving threat landscapes, and the need for secure IaC pipeline validation. Additionally, false positives remain a concern when threshold tuning isn't aligned with dynamic workload behaviors.

Future iterations should incorporate:
- Real-time SOAR feedback loop
- Explainable AI models for auditability
- Federated learning for cross-org threat intelligence without data exposure

The system's modularity allows it to be extended to other cloud platforms and container orchestration tools like Kubernetes, where network policies and pod-to-pod communication rules can be similarly validated. This portability makes it suitable for industries with multi-regulatory compliance requirements (e.g., financial services, healthcare, and government). One operational consideration is integration with DevSecOps pipelines for real-time policy injection during CI/CD. We are also exploring compatibility with service mesh architectures like Istio to enforce identity-aware policies.

VII-A. VISUAL DIAGRAMS AND ARCHITECTURE

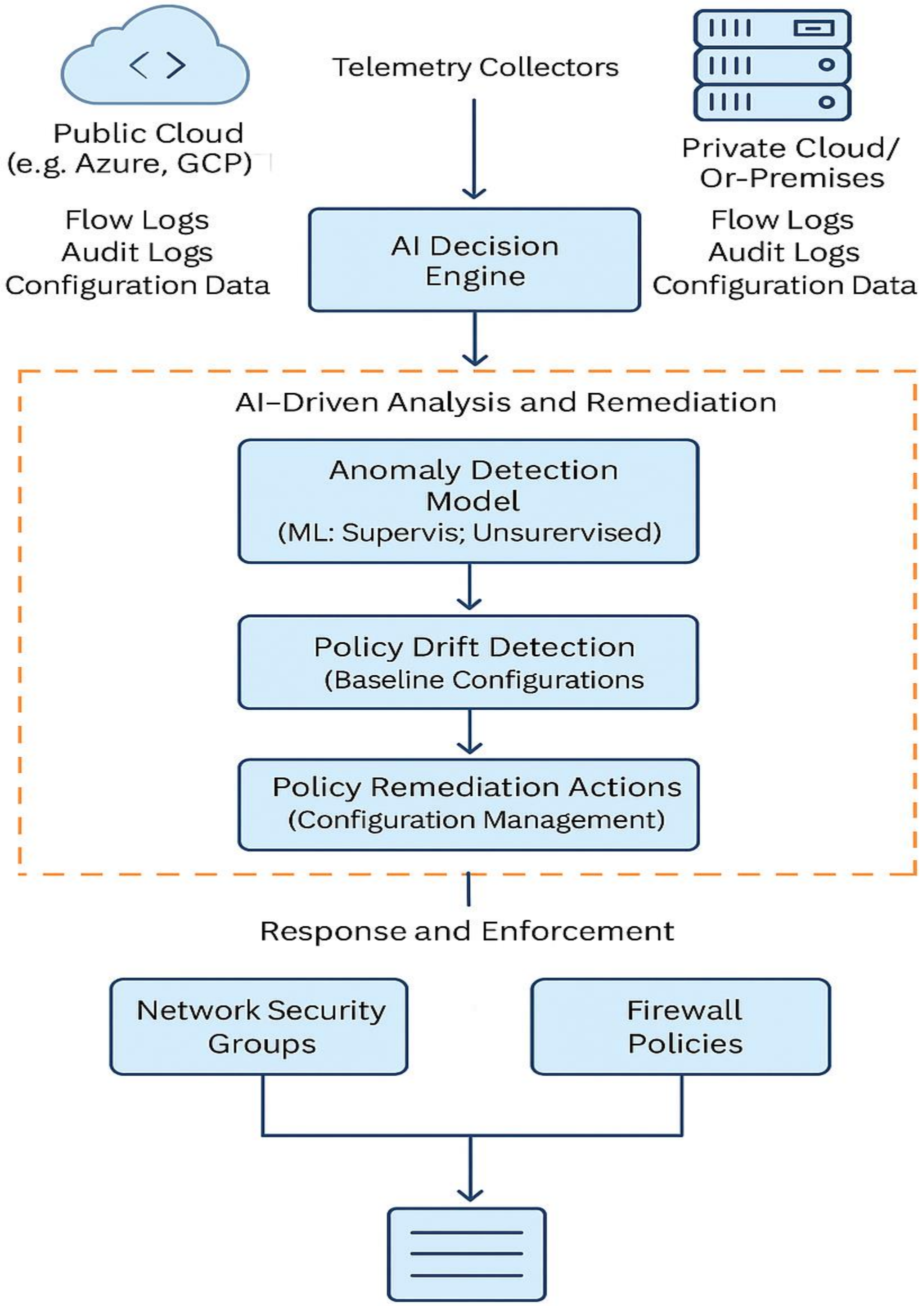

Figure 1: High-Level Architecture of the AI-Augmented Hybrid Cloud Security Framework

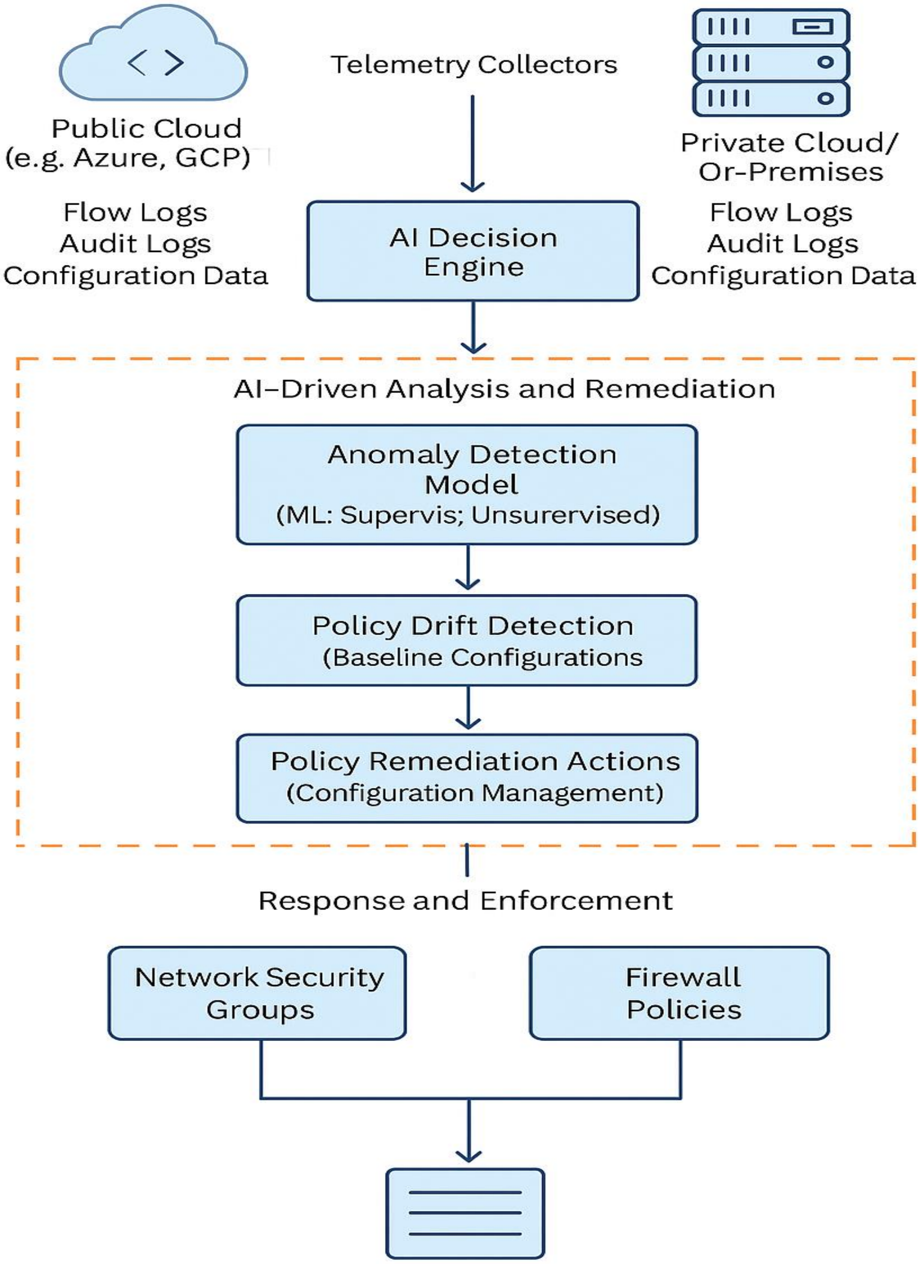

Figure 2: Flowchart of Detection and Remediation Workflow

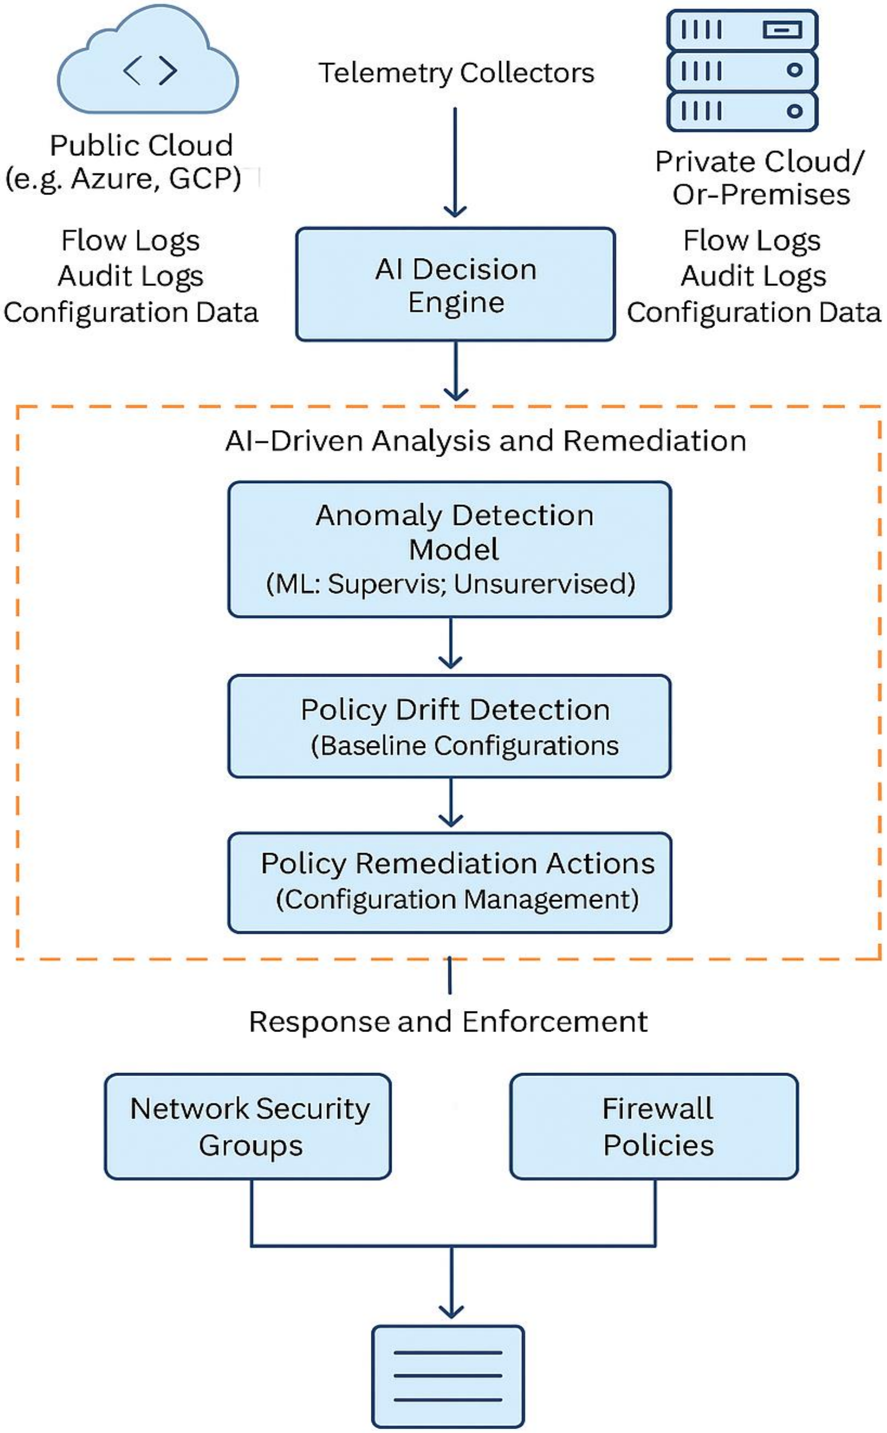

Figure 3: Model Pipeline for AI-Driven Anomaly Detection and Drift Remediation

VIII. CONCLUSION AND FUTURE WORK

This paper proposed and evaluated a full-stack AI-powered solution for threat detection and policy drift remediation in hybrid cloud environments. Combining telemetry aggregation, ML anomaly detection, policy reconciliation via IaC, and automated enforcement, the system enhances enterprise cloud posture resilience.

The hybrid testbed demonstrated high accuracy, fast response times, and reduced manual triage. Future work will focus on refining model transparency, integrating Kubernetes-native security enforcement, and extending support for multi-cloud and SaaS-based policy drift across more platforms like Oracle Cloud, Alibaba Cloud, and Salesforce.

REFERENCES

[1] Y. Zhang et al., 'Machine Learning-Based Network Anomaly Detection in Cloud Environments,' IEEE Trans. Cloud Computing, 2021.
[2] M. Tan et al., 'Deep Reinforcement Learning for Cybersecurity Policy Optimization,' in IEEE INFOCOM, 2022.
[3] Microsoft Azure Docs, 'Azure Network Watcher Flow Logs', 2023.
[4] Google Cloud, 'Security Command Center Documentation', 2023.
[5] A. Sharma et al., 'Policy Drift Detection for Cloud Infrastructure,' in IEEE CloudCom, 2020.
[6] Caldera Framework – MITRE ATT&CK, 'Adversary Emulation', 2023.
[7] Terraform Cloud Docs, 'Policy as Code with Sentinel', HashiCorp, 2023.
[8] Fluent Bit, 'Cloud-Native Telemetry Collection Framework', CNCF.
[9] J. Lee et al., 'Automated IaC Compliance and Security Posture Management,' in IEEE ICC, 2021.
[10] Palo Alto Networks, 'Cloud Security Automation Using ML', Whitepaper, 2022.