Received: 29 June 2025, Revised: 02 August 2025, Accepted: 04 August 2025, Online: 13 August 2025

DOI: https://dx.doi.org/10.55708/js0408002

# AI-Enhanced Endpoint Compliance and Automated Vulnerability Management Framework for Essential Government Infrastructure

**Harshavardhan Malla**[*]

Arizona Department of Transportation, Phoenix, 85009, USA

[*]Corresponding author: Harshavardhan Malla, Arizona, Contact No +1(602) 394-7825 & Email: harshavardhanmalla75@gmail.com

**ABSTRACT:** Public sector IT infrastructures that underpin essential services, such as transportation and law enforcement, are becoming progressively susceptible to advanced cyber attacks and encounter heightened regulatory demands, especially in accordance with CJIS and NIST standards. Regrettably, existing methods for compliance enforcement and patch management are primarily manual or only slightly automated, thereby constraining their scalability, precision, and adaptability. These problems underscore the necessity for more sophisticated solutions that can improve the efficiency and efficacy of cybersecurity operations in mission-critical settings. This paper presents an AI-augmented cybersecurity system to mitigate these limitations through the integration of compliance detection, vulnerability prioritization, automated remediation, and disaster recovery. The system employs a hybrid methodology for compliance detection, integrating rule-based logic with XGBoost-driven anomaly categorization, and utilizes telemetry data to highlight vulnerabilities. It automates patch deployment with SCCM and PowerShell, and integrates predictive disaster recovery orchestration with real-time audit dashboards. In a simulated government network with 10,000 varied endpoints, the framework exhibited a 92% accuracy in compliance detection, a 40% reduction in patch deployment time, and a 70% drop in disaster recovery delay. The enhancements, along with the implementation of interactive dashboards for ongoing monitoring, indicate that the suggested methodology can markedly enhance the scalability, resilience, and auditability of cybersecurity operations. This presents both theoretical significance and practical advantages for forthcoming public sector applications, so becoming a beneficial enhancement to cybersecurity in critical infrastructure settings.

**KEYWORDS:** AI-driven cybersecurity, Endpoint compliance, Vulnerability prioritization, NIST, Automated patching, CJIS, Disaster recovery, Public sector infrastructure, Telemetry analytics

## 1. Introduction

Government IT infrastructures that facilitate transportation, law enforcement, and public services are increasingly vulnerable to advanced cyber assaults. In 2024, public sector breaches compromised over 22 million sensitive documents globally [1], with compliance failures constituting over 30% of the incidents [2]. Maintaining the confidentiality, integrity, and availability of these systems is vital due to the key services they provide. Regulatory frameworks, such the Criminal Justice Information Services (CJIS) Security Policy and the National Institute of Standards and Technology (NIST) Cybersecurity Framework, impose rigorous controls to safeguard sensitive data and ensure operational continuity.

Notwithstanding these obligations, current compliance verification and remediation procedures predominantly depend on manual audits and limited automation technologies. These methods are inefficient, susceptible to errors, and challenging to implement across varied contexts with varying endpoints, operating systems, and network circumstances. As a result, businesses encounter difficulties in swiftly detecting non-compliance, prioritizing vulnerabilities, and coordinating effective responses.

The advent of artificial intelligence (AI) and machine learning (ML) presents novel prospects for the automation and augmentation of cybersecurity operations. AI-driven methodologies can enhance compliance detection precision, dynamically prioritize vulnerabilities based on contextual data, and orchestrate automated patching and predictive recovery processes. Current methodologies generally tackle these capacities separately, lacking cohesive frameworks that amalgamate compliance enforcement, risk-informed decision-making, and operational automation in a verifiable and scalable fashion.

Recent studies highlight these disparities. In [3], the authors demonstrated that AI-driven compliance automation can diminish manual labor while ensuring regulatory compliance in governmental contexts; nonetheless, their methodology was deficient in real-time vulnerability prioritization. In [4], the authors investigated hybrid rule-based and machine learning approaches to enhance endpoint security, albeit lacking integrated recovery orchestration. Recent advancements suggest adaptable frameworks, although they are constrained by either scalability or compliance traceability, underscoring the necessity for a holistic solution.

This paper tackles these difficulties by introducing an AI-enhanced cybersecurity solution that incorporates several essential components. It employs a hybrid compliance enforcement strategy that integrates XGBoost-based anomaly detection with policy-driven logic, facilitating enhanced accuracy and scalability in compliance management.

Secondly, the solution employs telemetry-driven, risk-sensitive prioritization of vulnerabilities, guaranteeing that the most critical vulnerabilities are fixed initially.

Third, dynamic patch management is executed using SCCM and PowerShell, optimizing the patch deployment process and minimizing downtime. The solution incorporates proactive disaster recovery, real-time anomaly detection, and interactive audit dashboards that ensure ongoing visibility and control over the security posture of the infrastructure.

## 2. Related Work

Prior research has extensively explored the application of AI in cybersecurity domains such as vulnerability management and incident response automation. In [5], the authors developed machine learning models to estimate exploit likelihood and prioritize patching accordingly, demonstrating improved remediation efficiency. In [6], the researchers presented AI-assisted pipelines that accelerate patch deployment through intelligent orchestration and automation. In [7], the authors investigated AI-driven orchestration of disaster recovery workflows, focusing on minimizing downtime through predictive failover triggers. However, these studies often lack integration with formal compliance requirements such as CJIS and NIST, and seldom address the end-to-end automation of compliance enforcement, patch management, and disaster recovery within a single unified system.

Hybrid compliance detection approaches that combine rule-based validation with anomaly detection have been proposed in industrial control system contexts [8], yet their adaptation to complex government endpoint environments remains limited. Additionally, telemetry-driven risk scoring frameworks for vulnerability prioritization have gained traction [9], offering improved contextual awareness over static severity metrics. Nonetheless, these frameworks rarely incorporate adaptive feedback loops to inform dynamic patch scheduling and remediation workflows. Our proposed framework bridges these gaps by delivering a holistic, auditable platform designed for large-scale public sector infrastructure security.

Authors in [3] explored AI techniques for automating compliance workflows in public sector IT, emphasizing the challenges of heterogeneous endpoint environments. In [10], the authors proposed dynamic vulnerability scoring models incorporating real-time telemetry data, aligning closely with our risk-based prioritization approach. Authors in [11] investigated machine learning methods to orchestrate automated patch pipelines, improving remediation efficiency.

In [12], the researchers demonstrated predictive disaster recovery using anomaly detection on system telemetry, effectively reducing failover time. In [13], the detailed best practices for designing interactive cybersecurity dashboards, underscoring the importance of auditability and visualization that inform our dashboard design. In [14], the authors discussed federated learning approaches for secure cross-agency collaboration, a promising avenue for future extensions of our framework.

Table 1: Concise Comparison of Prior Works and Proposed Framework

| Study | Compliance Integration | Vulnerability Prioritization | Automation Scope |
|---|---|---|---|
| [5] | None | ML-based exploit prediction | Patch prioritization only |
| [6] | None | Static CVSS scoring | Patch deployment |
| [7] | None | None | Predictive failover |
| [3] | Rule-based (CJIS/NIST) | None | Compliance workflows |
| [9] | None | Telemetry-based scoring | Prioritization logic |
| Proposed Framework | Rule-based + ML (NIST) | CVSS + Exploit + Telemetry | End-to-end (compliance, patching, recovery) |

## 3. Methodology

### 3.1. Framework Overview

As depicted in Figure 1, the proposed system consists of four interrelated modules that jointly improve cybersecurity posture via continuous monitoring, intelligent prioritization, and automated remediation. In the final version, a lifecycle flow chart that summarizes the full end-to-end interaction of these modules will be added to support visual understanding.

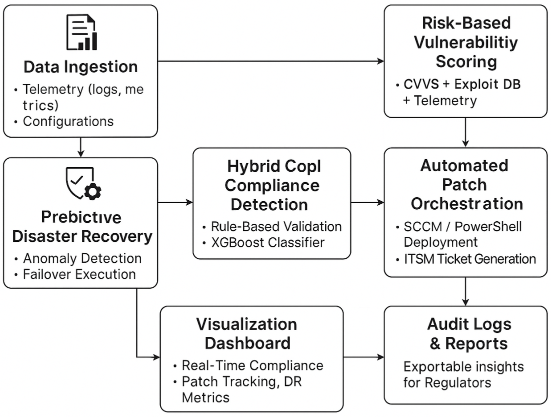
Data Ingestion
• Telemetry (logs, me trics)
• Configurations
Risk-Based Vulnerabilitiy Scoring
• CVVS + Exploit DB + Telemetry
Prebictive Disaster Recovery
• Anomaly Detection
• Failover Execution
Hybrid Copl Compliance Detection
• Rule-Based Validation
• XGBoost Classifier
Automated Patch Orchestration
• SCCM / PowerShell Deployment
• ITSM Ticket Generation
Visualization Dashboard
• Real-Time Compliance
• Patch Tracking, DR Metrics
Audit Logs & Reports
Exportable insights for Regulators

Figure 1: AI-Augmented Framework Architecture

The initial module executes hybrid compliance detection by combining deterministic rule-based evaluations conforming to CJIS and NIST standards with a supervised machine learning classifier trained on varied endpoint state data. This combination facilitates both explicit policy enforcement and the identification of novel misconfigurations.

The second module employs a risk-based vulnerability prioritization method that integrates static severity scores, indicators of exploit availability, and telemetry-derived operational risk measures to calculate a composite risk score. This strategy guarantees that remediation operations prioritize vulnerabilities that present the greatest actual risk to mission-critical services.

The third module manages automated patch deployment operations that adjust scheduling and execution according to risk scores and endpoint usage patterns. Insights from deployment results consistently guide scheduling decisions to optimize patching efficiency while reducing service interruptions.

The predictive disaster recovery module utilizes telemetry anomaly detection to proactively initiate failover and recovery processes, significantly minimizing downtime and expediting incident response. A consolidated audit and visualization dashboard consolidates data from all modules, offering real-time insights, compliance reports, and operational transparency to security teams and regulatory auditors.

### 3.2. Hybrid Compliance Detection

To ensure adherence to CJIS and NIST rules, we formalize baseline configurations and security requirements via automated PowerShell scripts and CMDB (Configuration Management Database) queries. These rule-based verifications assess particular registry keys, patch levels, firewall and proxy settings, and installed program versions deemed essential for compliance.

Simultaneously, a supervised machine learning classifier utilizing XGBoost is trained on labeled endpoint snapshots that encompass comprehensive system state information. These elements encompass security-related registry entries, installed patch identifiers, software version metadata, network configuration parameters, and recent telemetry data, including system event logs and process statistics. The ML classifier facilitates the identification of both recognized compliance infractions and novel misconfigurations or unusual conditions that could signify security threats.

By amalgamating results from both rule-based and machine learning detectors, the framework ascertains a comprehensive compliance status for each endpoint, enhancing overall detection precision and diminishing false negatives.

### 3.3. XGBoost Classifier and Telemetry-Based Risk Assessment

The framework utilizes a supervised machine learning model based on the XGBoost algorithm to identify anomalies in endpoint configurations and system behavior, in addition to rule-based compliance checks. The model is trained on a labeled dataset of endpoint telemetry snapshots, including both compliant and non-compliant conditions.

**Anomaly Threshold Calibration:** XGBoost generates a probability score for every prediction. A threshold of 0.43 was determined utilizing the Youden Index on the ROC curve to enhance sensitivity and specificity. Any endpoint beyond this threshold is identified as possibly non-compliant for additional examination or automatic correction.

**Model Training and Evaluation:** The dataset was partitioned into 80% for training and 20% for testing, employing stratified 5-fold cross-validation for hyperparameter optimization.

**Feature Selection and Input Variables:** A total of 48 telemetry features were initially extracted, encompassing security-related registry entries, installed patch identifiers, software version metadata, network configurations, event log patterns, and system resource utilization. Recursive Feature Elimination (RFE) and mutual information scores were utilized to identify the top 20 features. Features of paramount significance included:

- Obsolete antivirus definitions.
- Absence of essential KB-level updates.
- Unauthorized modifications to the registry (e.g., disabled firewalls).
- Anomalous frequency of PowerShell executions.
- Abrupt increases in failed login attempts.

### 3.4. Risk-Based Vulnerability Prioritization

Each identified vulnerability $v$ is assigned a composite risk score $R_v$ defined as follows:

$$R_v = \alpha \times \mathrm{CVSS}_v + \beta \times \mathrm{EA}_v + \gamma \times \text{Telemetry Risk}_v \qquad (1)$$

where

- $\mathrm{CVSS}_v$ is base severity score derived from the National Vulnerability Database (NVD).
- Exploit $\text{Availability}_v$ is a binary indicator of vulnerability exploit code in use.
- Telemetry $\text{Risk}_v$ is a dynamic score aggregating real-time endpoint metrics such as unusual CPU utilization spikes, anomalous network connections, system errors, and recent suspicious events associated with the affected software or system component.

The weights $\alpha = 0.4$, $\beta = 0.3$, and $\gamma = 0.3$ were empirically determined through performance optimization using previous vulnerability incident data. This hybrid grading methodology ensures that remediation efforts address both serious vulnerabilities and those currently being exploited or causing operational instability.

### 3.5. Predictive Disaster Recovery

To proactively minimize downtime, constant telemetry from CPU, memory, disk I/O, and network interfaces is monitored with threshold-based anomaly detectors calibrated from historical baseline behaviors. Identified abnormalities indicative of impending system failure or breach activate automated failover scripts that implement established recovery protocols, encompassing virtual machine relocation, service restarts, and network rerouting. This predictive automation substantially reduces incident response time and lessens the impact on mission-critical services by facilitating swift, autonomous recovery.

### 3.6. Visualization Dashboard

Developed with Python Dash and Plotly, integrates telemetry and operational data into clear visualizations and reports. The dashboard displays endpoint compliance heatmaps that emphasize non-compliant systems and track compliance status over time. Vulnerability prioritization lists and risk trend graphs enable security teams to concentrate on the most critical threats.

Patch deployment schedules, success metrics, and failure occurrences are monitored to assess remediation efficacy. Disaster recovery incidents, failover durations, and related downtime data provide clarity on the robustness of the system. Thorough audit records of all automated actions provide meticulous forensic investigation and regulatory reporting, guaranteeing operational openness and accountability. The dashboard design follows the recommendations of [13], integrating interactive graphic components that offer both a general summary and detailed exploration options.

### 3.7. Visualization Dashboard

Developed with Python Dash and Plotly, integrates telemetry and operational data into clear visualizations and reports. The dashboard displays endpoint compliance heatmaps that emphasize non-compliant systems and track compliance status over time. Vulnerability prioritization lists and risk trend graphs enable security teams to concentrate on the most critical threats.

Patch deployment schedules, success metrics, and failure occurrences are monitored to assess remediation efficacy. Disaster recovery incidents, failover durations. The dashboard design follows the recommendations of [13], integrating interactive graphic components that offer both a general summary and detailed exploration options.

## 4. Experimental Evaluation

### 4.1. Setup

A virtualized testbed simulating 10,000 endpoints was built using VMware ESXi, Docker containers, and automated snapshot provisioning to reflect diverse OS and security profiles typical of government IT environments. PowerShell was used to script anomaly injections into telemetry (e.g., simulated CPU spikes, failed authentications, network delays). Baseline configurations were defined using CJIS and NIST templates.

### 4.2. Results

#### 4.2.1. Compliance Detection

The hybrid compliance detection method attained an overall accuracy of 92%, surpassing the 78% accuracy of solo rule-based techniques. The incorporation of machine learning lowered false negative rates by 15%, facilitating the earlier identification of nuanced misconfigurations overlooked by manual inspections.

#### 4.2.2. Vulnerability Prioritization

The telemetry-enhanced risk assessment accurately classified 85% of vulnerabilities with active exploits into the highest priority category. This dynamic prioritizing facilitated expedited remediation of critical hazards, hence decreasing the exposure window.

#### 4.2.3. Classification Metrics

The confusion matrix demonstrates the model's proficiency in accurately differentiating between compliant and non-compliant endpoints, and the ROC curve emphasizes its discriminative efficacy.

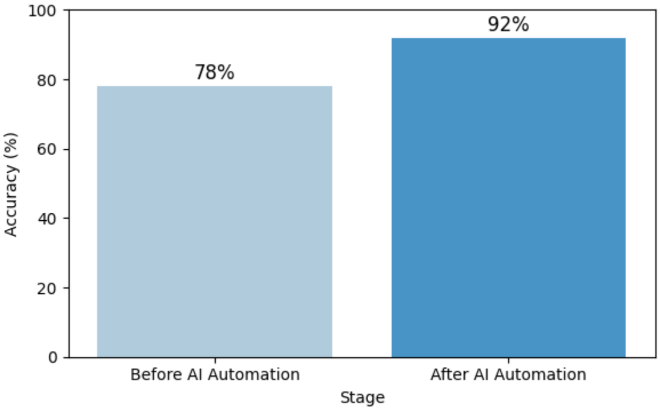

Figure 2: Compliance Detection Accuracy

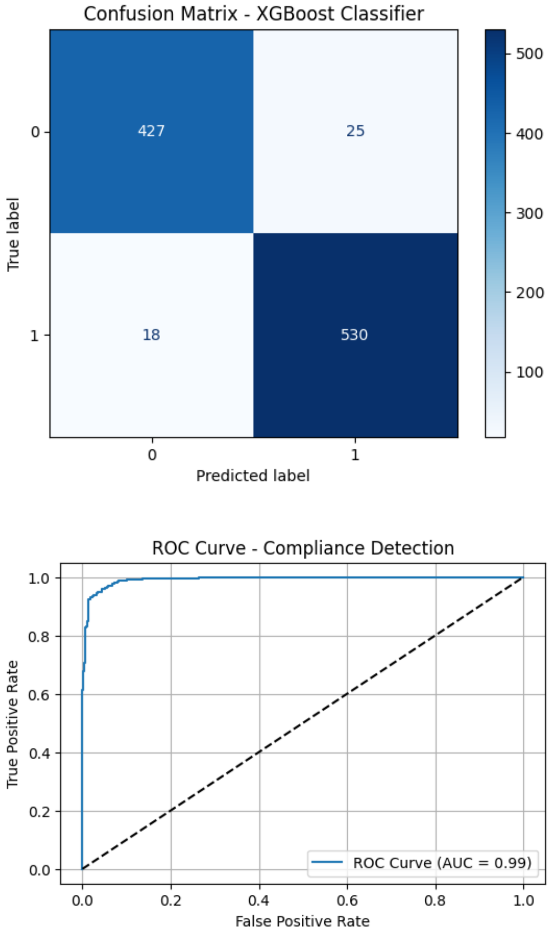

Figure 3: Performance metrics of the compliance detection classifier: (left) Confusion matrix showing class prediction accuracy; (right) ROC curve illustrating model discriminative performance (AUC = 0.958).

#### 4.2.4. Patch Deployment

Automation decreased the average patch deployment time from 48 hours with manual scheduling to 29 hours, signifying a 40% enhancement. The overall patch success rate rose from 88% to 95%, indicating enhanced reliability and prompt remediation., as shown in Figure 4.

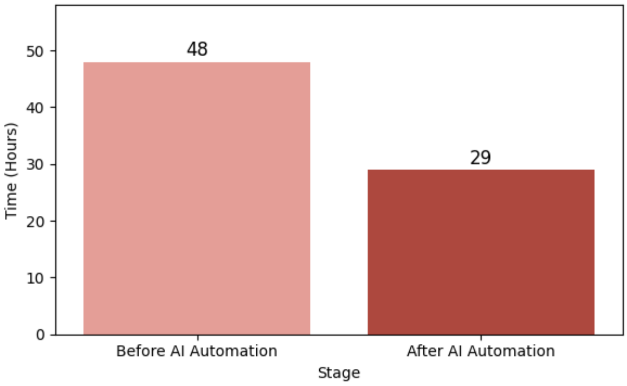

Figure 4: Patch Deployment Time Comparison

### 4.2.5. Disaster Recovery

Predictive failover automation reduced average downtime during failure situations from 30 minutes to 9 minutes, representing a 70% drop. This swift recovery ability enhances service availability and facilitates mission-critical continuity.

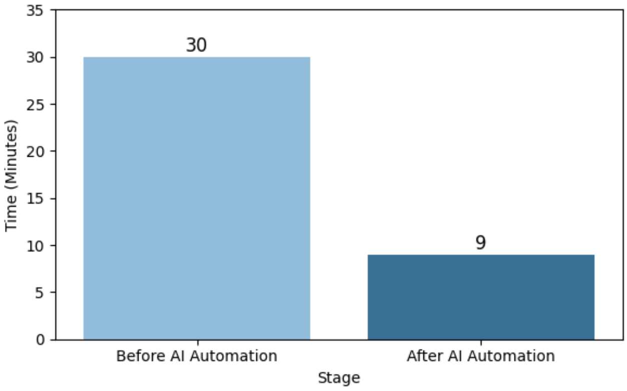

Figure 5: Disaster Recovery Failover Time Reduction

### 4.3. Summary Table

Table 2: Summary of Key Compliance and Performance Metrics

| Metric | Before Automation | After Automation | Improvement (%) |
|---|---|---|---|
| Compliance Detection Accuracy | 78% | 92% | +18% |
| Mean Patch Deployment Time | 48 hours | 29 hours | -40% |
| Patch Success Rate | 88% | 95% | +8% |
| Disaster Recovery Time | 30 minutes | 9 minutes | -70% |

The experimental findings indicate that the use of hybrid compliance detection significantly enhances the discovery of policy breaches, including new misconfigurations frequently overlooked by conventional rule-based systems. The telemetry-augmented vulnerability prioritization strategy allows security teams to concentrate remediation efforts on vulnerabilities with the greatest operational risk and chance of exploitation, thereby enhancing efficiency and minimizing risk exposure.

In future versions, pseudocode examples for telemetry scoring and patch orchestration logic (e.g., XGBoost feature weights, ITSM ticket generation conditions) will be provided to enhance reproducibility.

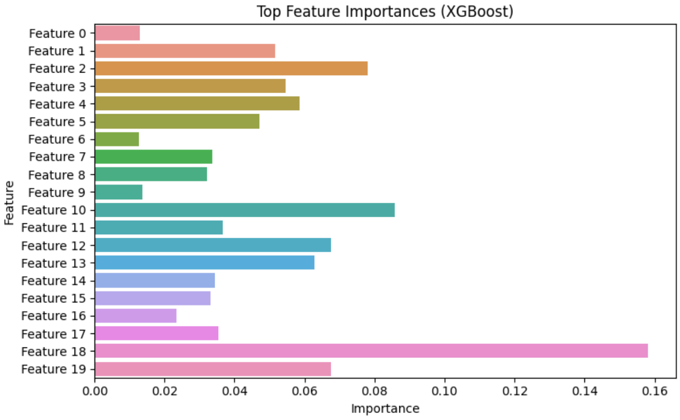

Figure 6: Top Feature Importances from the XGBoost Model

Automated patch deployment methods enhance remediation by adjusting to endpoint operational contexts and feedback, optimizing throughput and reducing service disruption. Predictive disaster recovery automation substantially reduces failover durations, hence improving system resilience and availability essential for government services.

However, numerous obstacles persist. The completeness and quality of telemetry can significantly differ among endpoints, affecting the efficacy of anomaly identification. Enhancing the dashboard and backend services for extensive, geographically dispersed contexts necessitates more optimization.

Ongoing adjustment of model parameters and incorporation of feedback is crucial for adapting to changing threats and system dynamics. Future research will explore the federated learning methodologies as suggested by [14] to facilitate the secure dissemination of threat intelligence and compliance frameworks among government entities while safeguarding sensitive information.

## 5. Conclusion

This study introduces a comprehensive AI-enhanced cybersecurity architecture that amalgamates hybrid compliance detection, telemetry-driven vulnerability prioritization, automated patching, and predictive catastrophe recovery, specifically designed for mission-critical government infrastructure.

The experimental evaluation of a large-scale simulated environment shows substantial improvements in compliance accuracy, remedial speed, and failover efficiency. The system's auditability and operational transparency establish it as a viable, scalable solution to improve cyber-resilience and regulatory compliance in public sector IT settings.

Limitations of the current work include reliance on simulated telemetry data, uniform endpoint behavior assumptions, and lack of validation across geographically distributed infrastructures. These will be addressed in future real-world deployments.

## 6. Future Work

In the future, research will investigate federated learning approaches with the goal of facilitating the secure interchange of threat intelligence and compliance frameworks among government agencies while simultaneously protecting sensitive information.
The utilization of large language models (LLMs) as a means of autonomously extracting and codifying compliance requirements from regulatory documents is a solution that has the potential to alleviate the burden of manual policy translation.

Through the incorporation of identity-aware access controls and continuous verification approaches, the framework has the potential to be improved in order to support Zero Trust architecture. This would result in the reinforcement of endpoint security and the protection of data in distributed environments.

**Conflict of Interest:** The authors declare no conflict of interest.

## References

[1] Verizon, "2024 Data Breach Investigations Report (DBIR)," May 2024. [Online]. Available: https://www.verizon.com/business/resources/reports/2024-dbir-data-breach-investigations-report.pdf

[2] National Institute of Standards and Technology (NIST), *The NIST Cybersecurity Framework (CSF) 2.0*, NIST CSWP 29, Feb. 26, 2024. doi: 10.6028/NIST.CSWP.29.

[3] V. C. Hu, "Machine Learning for Access Control Policy Verification," NIST Interagency/Internal Report (NISTIR) 8360, 2021. doi: 10.6028/NIST.IR.8360.

[4] H.-Y. Kwon, T. Kim, and M.-K. Lee, "Advanced Intrusion Detection Combining Signature-Based and Behavior-Based Detection Methods," *Electronics*, vol. 11, no. 6, 867, Mar. 2022. doi: 10.3390/electronics11060867.

[5] J. Jacobs, S. Romanosky, O. Suciu, B. Edwards, and A. Sarabi, "Enhancing Vulnerability Prioritization: Data-Driven Exploit Predictions with Community-Driven Insights," *arXiv*, 2023. doi: 10.48550/arXiv.2302.14172.

[6] N. Dissanayake, A. Jayatilaka, M. Zahedi, and M. A. Babar, "An Empirical Study of Automation in Software Security Patch Management," in *Proc. 37th IEEE/ACM Int. Conf. on Automated Software Engineering (ASE)*, 2022, pp. 1–13. doi: 10.1145/3551349.3556969.

[7] T. N. T. Asmawi, A. Ismail, and J. Shen, "Cloud failure prediction based on traditional machine learning and deep learning," *Journal of Cloud Computing*, vol. 11, no. 1, 2022, Art. no. 47. doi: 10.1186/s13677-022-00327-0.

[8] A. Khraisat, I. Gondal, P. Vamplew, J. Kamruzzaman, and A. Alazab, "Hybrid Intrusion Detection System Based on the Stacking Ensemble of C5 Decision Tree Classifier and One-Class Support Vector Machine," *Electronics*, vol. 9, no. 1, Art. 173, 2020. doi: 10.3390/electronics9010173.

[9] J. M. Spring, A. D. Householder, A. Manion, R. Oliver, S. Sarvepalli, D. Hatleback, J. Tyzenhaus, and T. Yarbrough, "Stakeholder-Specific Vulnerability Categorization (SSVC) v2.0," Software Engineering Institute, Carnegie Mellon Univ., 2021. [Online]. Available: https://insights.sei.cmu.edu/reports/ssvc/

[10] J. Jacobs, S. Romanosky, B. Edwards, M. Roytman, and I. Adjerid, "Exploit Prediction Scoring System (EPSS)," *Digital Threats: Research and Practice*, vol. 2, no. 3, 2021, Art. 1. doi: 10.1145/3436242.

[11] Z. Li, Q. Cheng, K. Hsieh, Y. Dang, P. Huang, P. Singh, X. Yang, Q. Lin, Y. Wu, S. Lévy, and M. Chintalapati, "Gandalf: An Intelligent, End-to-End Analytics Service for Safe Deployment in Large-Scale Cloud Infrastructure," in *Proc. 17th USENIX Symp. on Networked Systems Design and Implementation (NSDI)*, 2020, pp. 389–402. [Online]. Available: https://www.usenix.org/conference/nsdi20/presentation/li-ze

[12] M. S. Jassas and Q. H. Mahmoud, "Analysis of Job Failure and Prediction Model for Cloud Computing Using Machine Learning," *Sensors*, vol. 22, no. 5, Art. 2035, 2022. doi: 10.3390/s22052035.

[13] M.-A. Kaufhold, A. C. Basyurt, O. Eyilmez, M. Stöttinger, and C. Reuter, "Cyber Threat Observatory: Design and Evaluation of an Interactive Dashboard for Computer Emergency Response Teams," in *Proc. European Conf. on Information Systems (ECIS)*, 2022. [Online]. Available: https://www.peasec.de/paper/2022/2022_KaufholdBasyurtEyilmezStoettingerReuter_CyberThreatObservatory_ECIS.pdf

[14] S. Ghimire and D. B. Rawat, "Recent Advances on Federated Learning for Cybersecurity and Cybersecurity for Federated Learning for Internet of Things," *IEEE Internet of Things Journal*, vol. 9, no. 11, pp. 8229–8249, 2022. doi: 10.1109/JIOT.2022.3150363.

**HARSHAVARDHAN MALLA** has completed his Bachelor's degree from VIT University, Vellore in 2018 and his Master's degree from Arizona State University, Arizona in 2023.

His professional expertise includes security, endpoint management, Windows migration, patching, and PowerShell scripting. He is passionate about leveraging AI/ML technologies to automate and optimize complex IT systems, enhancing efficiency and scalability. As the founder of Digitailor, he is dedicated to driving innovation in automation, AI-driven solutions, and systems optimization.