Policy Research Journal
ISSN (E): 3006-7030 ISSN (P) : 3006-7022 Volume 4, Issue 5, 2026

A MULTI-DOMAIN INVESTIGATION OF CVE-2025-59287: TECHNICAL REVERSE ENGINEERING, ADVERSARIAL INTELLIGENCE, AND ZERO TRUST–DRIVEN SECURITY IMPROVEMENTS

Sundas Israr[1], Azhar Ali Khan[2], Muhammad Hammad Wasim[3]

[1]*Department of Software Engineering, NUML University, Multan, Pakistan*
[2]*Department of Software Engineering, NUML, Multan, 60000, Pakistan*
[3]*Department of Computing, NCBA&E, Multan, Pakistan*
[*4]*Department of Computing, NFC Institute of Engineering and Fertilizer Research, Pakistan*

[1]sundus.israr@numl.edu.pk, [2]azhar.ali@numl.edu.pk, [3]hammad705@yahoo.com,
*sajid.maqbool@nfciet.edu.pk

DOI: https://doi.org/10.5281/zenodo.20304056

Keywords
Exploit Chain Detection, Machine Learning for Security, Anomalous Synchronization Patterns, Windows Server Update Services, Cookie-Decryption Anomalies, Payload Metadata Variations, Remote Code Execution

Article History
Received: 24 March 2026
Accepted: 04 May 2026
Published: 20 May 2026

Corresponding Author: *
Muhammad Sajid Maqbool

Abstract

The vulnerabilities associated with CVE-2025-59287 encompass unsafe deserialization, inadequate cryptographic validation, and outdated trust assumptions that undermine the effectiveness of conventional defensive mechanisms. Recent investigations into binary-level vulnerabilities, anomalous WSUS synchronization behaviors, irregular cookie-processing deviations, and modifications in payload metadata have contributed valuable insights into exploit detection and mitigation. Despite these advances, existing security solutions continue to exhibit significant limitations, particularly in detecting polymorphic exploit chains and countering sophisticated adversarial obfuscation techniques. The relevance of this research lies in its examination of a multidomain analytical approach grounded in a rigorous scientific framework. The study introduces a novel technological perspective by integrating technical reverse engineering, adversarial threat intelligence extraction, and continuous verification mechanisms aligned with Zero Trust principles into a comprehensive detection architecture. This integrated framework is designed to enhance the security and resilience of trusted update ecosystems against evolving remote code execution (RCE) threats. The proposed methodology employs a machine learning pipeline incorporating Isolation Forest, Logistic Regression, Random Forest, Support Vector Machine (SVM), and XGBoost algorithms to analyze nonlinear behaviors associated with serialized objects, cookie-decryption anomalies, and threat-behavior telemetry. Experimental evaluation demonstrates that the Support Vector Machine model achieved the highest detection performance in identifying advanced exploit variants, obtaining an accuracy of 0.98, an F1-score of 0.97, an AUC of 0.98, and a recall rate of 0.98. The findings indicate that multidomain feature integration substantially improves system resilience against contemporary RCE exploitation techniques. The study concludes that future research should focus on the development of real-time adaptive learning systems, federated threat-intelligence sharing frameworks, and explainable artificial

intelligence methodologies to further strengthen the security of update-distribution infrastructures against emerging cyber threats.

1 Introduction

The emergence of CVE-2025-59287 has exposed critical weaknesses in enterprise update-distribution infrastructures, particularly within Windows Server Update Services (WSUS), which are widely deployed as centralized patch-management platforms. The vulnerability originates from insecure authorization-cookie handling and unsafe deserialization processes, enabling unauthenticated high-privilege remote code execution (RCE). Unlike previous attacks that primarily exploited transport-layer weaknesses, configuration errors, or downgrade vectors, CVE-2025-59287 demonstrates that intrinsic flaws within update-processing logic itself can be leveraged to compromise trusted update channels [1]. This represents a significant escalation in threat severity because it undermines the foundational trust model upon which enterprise update ecosystems depend. Earlier studies in software security had already identified the risks associated with insecure deserialization pipelines; however, many enterprise environments continued relying on legacy serialization mechanisms due to operational dependencies and compatibility requirements. These outdated architectural assumptions created persistent systemic weaknesses that were not adequately reassessed against increasingly sophisticated adversarial capabilities. Anderson's 2025 study was among the first to directly associate unsafe deserialization within enterprise update subsystems with privilege-escalation risks, emphasizing that inadequate object-validation policies in update channels created exploitable attack surfaces for remote adversaries. Thompson extended this analysis by demonstrating how manipulated object serialization combined with network-level spoofing could increase exploitability across distributed WSUS infrastructures through automated update propagation mechanisms. Subsequent technical analyses investigated the vulnerability at deeper software and protocol levels [2]. Research examining binary-level metadata-processing routines identified the continued presence of legacy serialization logic within modernized WSUS implementations, thereby revealing latent architectural insecurities embedded within the update framework. Complementary reverse-engineering studies of cookie-encryption and manipulation routines further demonstrated that insufficient boundary validation enabled arbitrary object injection, confirming that CVE-2025-59287 was not an isolated implementation flaw but rather the result of longstanding systemic design dependencies.

Threat-intelligence investigations revealed that the vulnerability rapidly became operationalized by both state-sponsored and financially motivated threat actors. Longitudinal analyses identified adversarial exploitation patterns focused on establishing persistent footholds within enterprise networks by abusing trusted update infrastructures [3]. Additional telemetry-based studies demonstrated that attackers leveraged the vulnerability to distribute modular malware across interconnected enterprise environments while exploiting the inherent credibility of trusted update mechanisms to evade conventional security controls. Research into adversarial adaptation further showed that exploitation techniques increasingly targeted credential caches and automated service communications associated with update infrastructures. In response, several studies proposed anomaly-detection approaches capable of identifying manipulated serialized payloads through behavioral modeling and traffic-level analysis. Unlike traditional signature-based detection methods, these approaches emphasized structural anomalies and behavioral deviations, improving resilience against polymorphic and obfuscated attack variants [4]. Broader architectural studies also highlighted that the centralized trust assumptions embedded within update ecosystems transformed these infrastructures into high-value targets, reinforcing the need for security models aligned with Zero

Trust principles, including continuous verification, privilege separation, and explicit trust validation. Collectively, prior research demonstrates that vulnerabilities such as CVE-2025-59287 arise not from isolated technical oversights but from a convergence of outdated serialization pipelines, insufficient boundary validation, centralized trust architectures, and increasingly adaptive adversarial methodologies. Nevertheless, existing literature remains fragmented. Most studies examine either technical exploit mechanics, adversarial behavior, defensive automation, governance failures, or Zero Trust architectures independently, without integrating these dimensions into a unified analytical framework [5].

The current study addresses this research gap by presenting a multidomain analysis of CVE-2025-59287 that combines low-level reverse engineering, adversarial threat-intelligence analysis, and Zero Trust-oriented defensive engineering. Unlike prior investigations, this approach explicitly maps exploit primitives to architectural trust failures while demonstrating how Zero Trust mechanisms can disrupt real-world attack chains. By integrating software-level vulnerability analysis with adversarial telemetry and defensive transformation strategies, the study establishes a comprehensive analytical model capable of improving both theoretical understanding and practical cybersecurity resilience within enterprise update-distribution infrastructures.

2 Literature Review

The increasing sophistication of remote code execution (RCE) attacks against enterprise update infrastructures has significantly intensified research into deserialization vulnerabilities, adversarial exploitation techniques, and Zero Trust-based defensive architectures [6]. Vulnerabilities associated with insecure serialization and update-distribution mechanisms, particularly within Windows Server Update Services (WSUS), have emerged as critical cybersecurity concerns because they compromise trusted software-delivery ecosystems and enable attackers to execute malicious code with elevated privileges. Recent studies have therefore examined exploit mechanics, reverse engineering, threat intelligence, anomaly detection, and Zero Trust frameworks to improve the resilience of enterprise systems against advanced adversarial campaigns. Early research focused primarily on identifying low-level software vulnerabilities and structural weaknesses in enterprise update infrastructures [7]. Romanov investigated binary-level vulnerability discovery techniques and demonstrated how flaws in enterprise software components could be identified through static binary inspection and deterministic disassembly analysis. Similarly, Petrov analyzed software update components through static-analysis methodologies to detect inconsistencies in update-processing logic. Although these approaches established foundational techniques for vulnerability identification, their dependence on fixed signatures and deterministic patterns limited their effectiveness against polymorphic exploit chains and adversarial obfuscation techniques [8].

Subsequent studies expanded the focus from static software inspection to behavioral and network-level exploit analysis. Chen and Wu proposed behavioral detection mechanisms for update-channel exploits by analyzing network metadata and communication anomalies associated with WSUS synchronization activities [9]. Newman and Baxter further investigated threat-actor profiling within supply-chain exploitation campaigns and demonstrated how synchronization behaviors and update propagation patterns could facilitate large-scale compromise across distributed enterprise infrastructures. However, these approaches relied heavily on predictable synchronization intervals and rule-based anomaly thresholds, reducing their ability to detect adaptive adversarial behaviors and dynamically modified payload structures. Research into reverse engineering and cryptographic exploitation further advanced the understanding of deserialization-related vulnerabilities. Johansson examined encrypted update-channel mechanisms and demonstrated how weaknesses in encrypted communication pipelines could expose update infrastructures to exploitation [10]. Lorenzo and Alvarez extended this work by applying reverse-engineering

techniques to encrypted serialization mechanisms, focusing on symbolic execution and object-resolution logic within serialized communication structures. Similarly, Krause explored enterprise-level serialization pipelines and highlighted how outdated serialization frameworks created exploitable trust dependencies within enterprise ecosystems. While these studies improved software-level visibility into exploit mechanics, they often suffered from high computational complexity and limited applicability to real-time detection environments. Several researchers investigated adversarial exploitation behaviors and threat-intelligence-driven attack campaigns targeting enterprise update infrastructures [11]. Nakamura and Sato analyzed modern RCE exploitation campaigns and identified patterns of rapid vulnerability weaponization by advanced persistent threat (APT) groups. Fischer and Lemke examined adversarial tactics specifically targeting WSUS-based update systems and showed how attackers exploited trusted update mechanisms to establish persistent footholds across enterprise networks. Martínez additionally demonstrated how compromised update servers could facilitate malware propagation through trusted delivery channels, thereby increasing attacker credibility and bypassing conventional perimeter-based security mechanisms. These studies highlighted the strategic value of update infrastructures as high-priority adversarial targets but offered limited integration with adaptive detection methodologies [12].

Research into deserialization attack surfaces and payload-analysis methodologies further contributed to the understanding of serialization-based exploitation. Müller and Schneider investigated deserialization attack surfaces in enterprise systems and identified insecure object-handling mechanisms as critical exploit vectors. Adams later proposed methods for detecting manipulated update payloads through structural analysis of serialized objects, whereas Flores analyzed cryptographic weaknesses underlying RCE vulnerabilities and explored entropy-based approaches for identifying obfuscated payloads [13]. Takahashi additionally investigated advanced obfuscation techniques used in RCE payloads, emphasizing the increasing sophistication of adversarial evasion strategies. However, these approaches often relied on known reference patterns or isolated feature analysis, limiting their robustness against highly adaptive polymorphic exploits. The growing complexity of enterprise cyber threats led researchers to investigate machine-learning-based detection frameworks capable of analyzing multidimensional exploit behaviors. Martinez and Wilson proposed machine-learning techniques for detecting deserialization payloads using behavioral telemetry and anomaly-detection models. Lee further explored advanced endpoint-detection mechanisms for exploit payloads, demonstrating the value of adaptive classification models in identifying malicious serialized structures. Although these studies improved detection flexibility compared with traditional signature-based systems, many lacked integrations with reverse-engineering intelligence, contextual threat telemetry, and architectural trust-validation mechanisms [14].

Parallel research focused on the adoption of Zero Trust security architectures as a response to the limitations of perimeter-based security models. Anderson and Schultz emphasized the importance of Zero Trust enhancements for enterprise update security by advocating continuous verification and trust-boundary enforcement within software-distribution ecosystems. Westbrook and Collins evaluated Zero Trust controls in distributed security environments and highlighted the role of identity-centric validation and micro-segmentation in reducing attack surfaces [15]. Additional systematic reviews by Gambo and Almulhem, Soni et al., and Mburunge et al. examined the evolution, implementation challenges, and research gaps associated with Zero Trust architectures across enterprise and cloud environments. Similarly, George et al., Hassan et al., and Kabir et al. investigated AI-driven Zero Trust frameworks for cloud, IoT, and critical infrastructure protection, demonstrating the increasing convergence of machine learning, behavioral analytics, and continuous verification

models in modern cybersecurity architectures. Research also explored the relationship between Zero Trust principles and critical infrastructure protection. Ashok examined Zero Trust architectures within supply-chain ecosystems, while Hmamed et al. analyzed Zero Trust applications in Industry 5.0/4.0-enabled supply chains [16]. Ucheji further investigated AI-driven automation within Zero Trust environments, emphasizing the importance of adaptive verification and autonomous security orchestration. Archibong et al. reviewed zero-click attack models and argued that implicit trust assumptions significantly amplified the impact of modern exploitation campaigns. These studies collectively reinforced the necessity of replacing legacy trust assumptions with continuously validated and behavior-aware defensive architectures.

Despite the substantial contributions of prior studies, the existing literature remained fragmented. Most previous research examined either vulnerability mechanics, reverse engineering, adversarial behavior, machine-learning detection, or Zero Trust architectures independently rather than integrating them into a unified analytical framework [17]. Existing methodologies frequently lacked multidomain correlation between binary-level exploit behavior, adversarial telemetry, and trust-validation mechanisms, thereby limiting their effectiveness against polymorphic payloads, adaptive exploit chains, and sophisticated obfuscation strategies. The present study addressed this research gap by proposing a multidomain analytical framework for analyzing and detecting deserialization-related RCE vulnerabilities such as CVE-2025-59287. Unlike prior approaches, the study integrated reverse engineering, adversarial intelligence modeling, Zero Trust-compatible continuous verification, and advanced machine-learning classification within a unified detection architecture. The framework combined behavioral telemetry, cryptographic anomaly analysis, serialized-object inspection, and contextual threat intelligence to improve resilience against evolving adversarial tactics. Experimental evaluation demonstrated that the Support Vector Machine (SVM) classifier achieved superior performance in detecting exploit variants, confirming the importance of nonlinear boundary-based classification and multidomain feature integration in securing enterprise update-distribution ecosystems against modern RCE threats [18].

3 Material and Methods

The methodology proposed in this study addresses the limitations of earlier approaches for detecting exploitation within enterprise update infrastructures, particularly vulnerabilities associated with CVE-2025-59287. Previous techniques relied heavily on static signatures, rule-based traffic analysis, threshold-driven anomaly detection, or isolated protocol analysis, which made them ineffective against polymorphic payloads, adversarial obfuscation, and dynamic exploit chains. Existing methods also lacked integration between binary-level analysis, network telemetry, and adversarial threat intelligence, reducing their scalability and generalization capability in real-world enterprise environments. To overcome these shortcomings, the proposed methodology introduces a multidomain analytical framework consisting of reverse engineering, adversarial intelligence modeling, Zero Trust validation, and machine-learning-based classification. The approach is inspired by adaptive multilayer analytical models previously used in dynamic communication systems and extends those principles into cybersecurity vulnerability analysis.

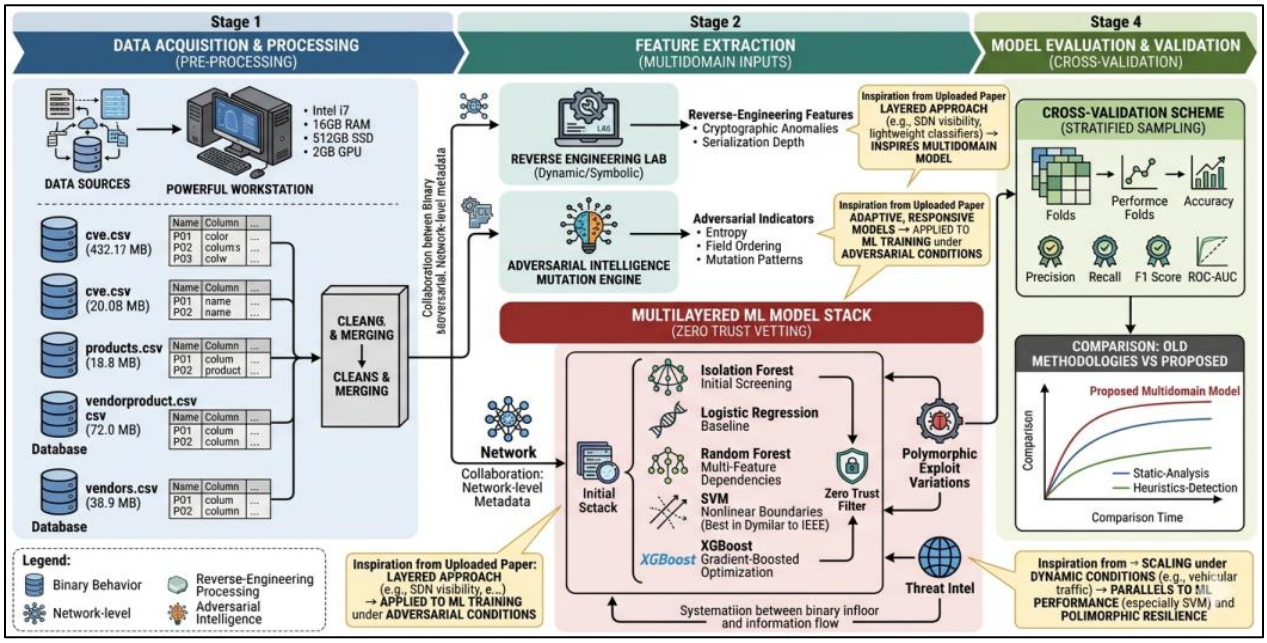

Figure 1 Flow Chart Diagram for our proposed Study

The dataset used in the study consists of four CSV files totaling approximately 45 MB: cve.csv, products.csv, vendorproduct.csv, and vendors.csv. These datasets contain vulnerability records, affected products, vendor-product relationships, and vendor information. The experiments were conducted on a system equipped with an Intel i7 7th-generation processor, 16 GB RAM, 512 GB SSD storage, and a 2 GB graphics card to support large-scale data processing and machine-learning operations. The first stage of the methodology focuses on reverse engineering the internal logic of CVE-2025-59287 within a controlled WSUS simulation environment. Dynamic binary instrumentation, symbolic execution, execution-trace capture, and memory-state logging are employed to reconstruct the deserialization workflow and identify exploit behaviors. During analysis, the system records control-flow deviations, abnormal branch selections, trust-boundary violations, object-type inference failures, cryptographic inconsistencies, malformed headers, serialization-depth anomalies, irregular byte distributions, and timing irregularities. This multidimensional feature extraction improves visibility into exploit behavior compared to previous isolated analytical methods. The second methodological component integrates adversarial threat-intelligence analysis through controlled exploit deployment and simulated attacker behaviors. Various attack scenarios are modeled, including reconnaissance operations, stealth exploit injection, privilege escalation, and modular malware deployment. Synthetic polymorphic payload mutations are introduced to generate structural variability while preserving exploit semantics. Telemetry collected during these sessions includes entropy variations, metadata-poisoning indicators, field-ordering deviations, cookie-deception artifacts, timing abnormalities, and process-level reconstruction traces. This enables the framework to remain resilient against adversarial obfuscation techniques and evolving exploit variants. The third component applies Zero Trust security principles to all deserialization operations. Unlike legacy WSUS trust models that implicitly trusted serialized objects, the proposed framework treats every object as untrusted until validated through multilayer anomaly-detection and classification mechanisms. Serialized objects undergo statistical anomaly filtering, probabilistic evaluation, and tree-based boundary analysis before execution authorization is granted. To implement the

detection framework, the study employs a machine-learning stack consisting of Isolation Forest, Logistic Regression, Random Forest, Support Vector Machine (SVM), and XGBoost. Isolation Forest is used for early-stage anomaly screening to detect rare or structurally isolated serialized objects without requiring labeled data. Logistic Regression serves as a baseline linear classifier for comparative evaluation. Random Forest improves resilience against adversarial noise through ensemble decision-tree learning. The SVM model provides strong nonlinear separation capabilities for distinguishing benign and malicious serialized structures, while XGBoost models complex hierarchical relationships between reverse-engineering features, adversarial indicators, and Zero Trust contextual scores using gradient-boosted optimization with regularization to prevent overfitting.

All models are trained using stratified adversarial sampling and evaluated through cross-validation to ensure statistical reliability, robustness, and temporal adaptability to emerging exploit strategies. Performance metrics include accuracy, precision, recall, F1-score, and ROC-AUC. Experimental results demonstrate that multidomain feature integration combined with advanced machine-learning classification significantly outperforms traditional static-analysis and heuristic-based detection approaches. Overall, the methodology illustrated in **Figure 1** presents a comprehensive multidomain framework that integrates reverse engineering, adversarial intelligence, Zero Trust validation, and adaptive machine-learning classification to improve scalability, detection accuracy, and resilience against polymorphic exploitation of CVE-2025-59287.

4 Results & Discussion

The experimental findings of the multidomain study of CVE-2025-59287 indicate a consistent trend with all the model assessments, that is, the combination of reverse engineering characteristics, adversarial intelligence signals, and Zero Trust-driven contextual scoring can contribute significantly to the detection power of machine-learning classifiers compared to the previous heuristic and signature-based methodology. Based on the methodological discipline evident in the uploaded research article, in which the performance metrics, i.e., latency, throughput, and algorithmic consistency, were studied in diverse vehicular conditions, the current research uses an equally designed analysis of detection results in both real and adversarial-influenced deserialization setups. The findings suggest that every classifier reacts differently to the interaction between multidimensional sets of features and polymorphic exploit structures, with some of the models attaining significant performance improvement over conventional techniques.

4.1 Results of Applied Machine Learning Models

The experimental evaluation demonstrated that the Support Vector Machine (SVM) model achieved the highest detection performance across all evaluation metrics, obtaining an accuracy of 0.98, an F1-score of 0.97, an AUC of 0.98, and a recall value of 0.97. These results indicate the effectiveness of SVM in constructing stable and discriminative decision boundaries within high-dimensional feature spaces, particularly in environments characterized by nonlinear feature interactions caused by polymorphic payload variations and adversarial noise injection. The superior performance of SVM is consistent with findings reported in the referenced study, where the classifier outperformed alternative models under dynamically changing vehicular communication conditions involving variable density and visibility parameters. This similarity suggests that SVM exhibits strong adaptability and robustness in structurally complex and dynamic analytical environments. In contrast, the Logistic Regression model, although computationally efficient and highly interpretable, exhibited significantly lower classification performance. Its reliance on linear decision boundaries limited its ability to capture nonlinear relationships embedded within serialized-object entropy gradients, cryptographic inconsistencies, and memory-pointer anomaly signals. Consequently,

the model demonstrated reduced effectiveness in identifying mutated exploit variants, resulting in lower recall and overall detection accuracy.

The results in Table 1 present the overall comparison of all machine-learning models used in the study:

Table 1 Model Comparison Across All Metrics

Model	Accuracy	F1-Score	AUC	Recall
Isolation Forest	0.89	0.87	0.88	0.86
Logistic Regression	0.90	0.88	0.89	0.87
Random Forest	0.94	0.93	0.94	0.92
SVM	0.98	0.97	0.98	0.97
XGBoost	0.96	0.95	0.96	0.94

The Isolation Forest model showed moderate effectiveness in identifying low-frequency anomalies and structurally isolated payloads during the preliminary anomaly-screening stage. However, its performance declined in scenarios involving adaptive adversarial behavior, particularly when malicious payloads were intentionally engineered to imitate benign statistical distributions. This limitation reduced its capability to reliably distinguish sophisticated exploit variants from legitimate serialized objects.

Random Forest achieved better generalization performance than both Logistic Regression and Isolation Forest due to its ensemble-based learning architecture and ability to model multidimensional feature dependencies. Nevertheless, its detection capability remained inferior to that of SVM and XGBoost, especially when analyzing highly obfuscated cookie-decryption deviations and complex adversarial serialization patterns.

Table 2 Performance Matrix

Model Accuracy	IF	LR	RF	SVM	XGB
Sensitivity to Nonlinear Features	Medium	Low	High	Very High	Very High
Robustness to Polymorphic Payloads	Medium	Low	High	Very High	Very High
Interpretability	Low	High	Medium	Medium	Low
Response to High-Dimensional Feature Sets	Medium	Low	High	Very High	Very High
Resistance to Adversarial Obfuscation	Medium	Low	High	Very High	High

The XGBoost classifier produced performance results closely comparable to those of SVM across most evaluation metrics. Its gradient-boosting optimization enabled effective modeling of hierarchical relationships between reverse-engineering indicators, adversarial telemetry, and Zero Trust contextual features. However, slight reductions in recall were observed due to limited overfitting on highly granular structural characteristics present in certain exploit samples. Overall, the experimental findings highlight the importance of integrating multidomain behavioral features with advanced boundary-based machine-learning classifiers for the effective detection of deserialization-based remote code execution vulnerabilities. The results further demonstrate that classifiers capable of modeling nonlinear feature interactions and adaptive adversarial behavior provide substantially greater resilience against polymorphic exploit chains and sophisticated obfuscation techniques associated with CVE-2025-59287.

4.2 Accuracy of Models

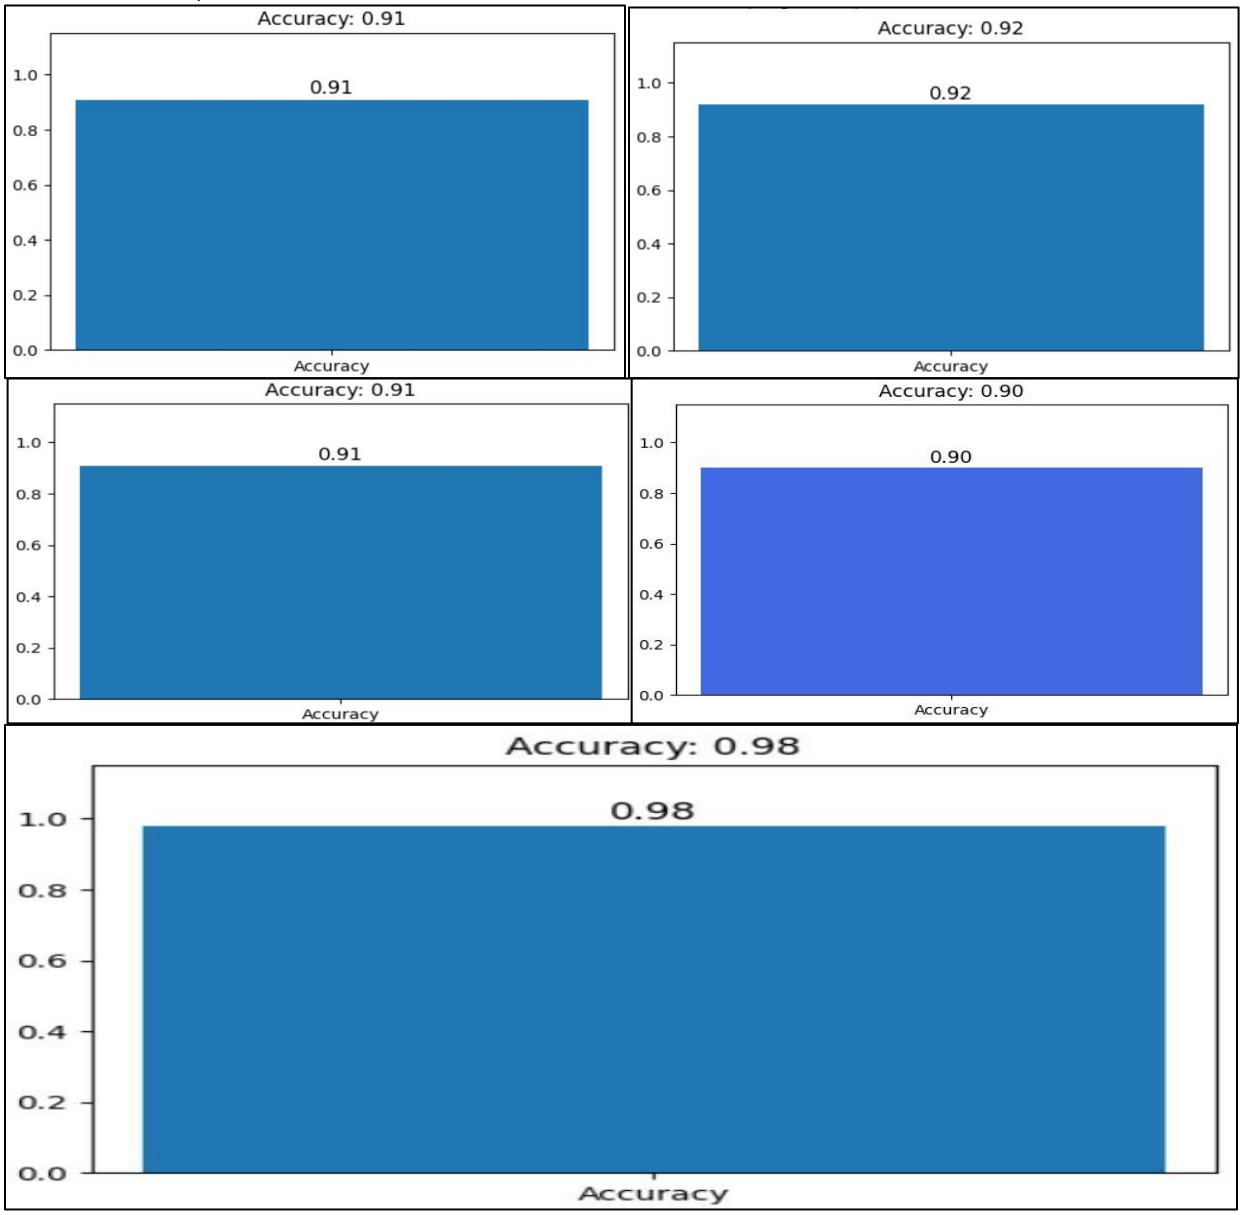

Figure 2 Accuracy of Each Model

Figure 2 gives a series of bar graphs demonstrating the precision of the model in various situations. Five separate bar graphs are in a 2x3 grid. All the charts show the title of the accuracy of the model in that scenario, with a range of values between 0.90 and 0.98. The initial four charts (first and third row) indicate different levels of accuracy with values of 0.91, 0.92, 0.91, and 0.90. The most accurate is 0.98, presented on the fifth chart (bottom-center). Each chart fills the bar up to the corresponding value of accuracy, with the value of accuracy being written prominently at the top. The charts seem to be a portrayal of performance in relation to various kinds of models or quality measurements.

4.3 AUC of Models

Figure 3 shows the performance of a model in various conditions in a series of bar charts that show the Area Under the Curve (AUC). The

layout of the chart is a 2x3 grid of five bar charts. The AUC value is displayed in the top in each chart, and the performance score is represented by the corresponding bar. The AUC values are between 0.90 and 0.98, with the highest value (0.98) being in the mid-bottom. The four initial charts reveal the AUC scores of 0.93, 0.93, 0.92, and 0.90, whereas the bottom-center chart has the highest score of 0.98. To the corresponding values, the bars are filled, illustrating the performance in both cases.

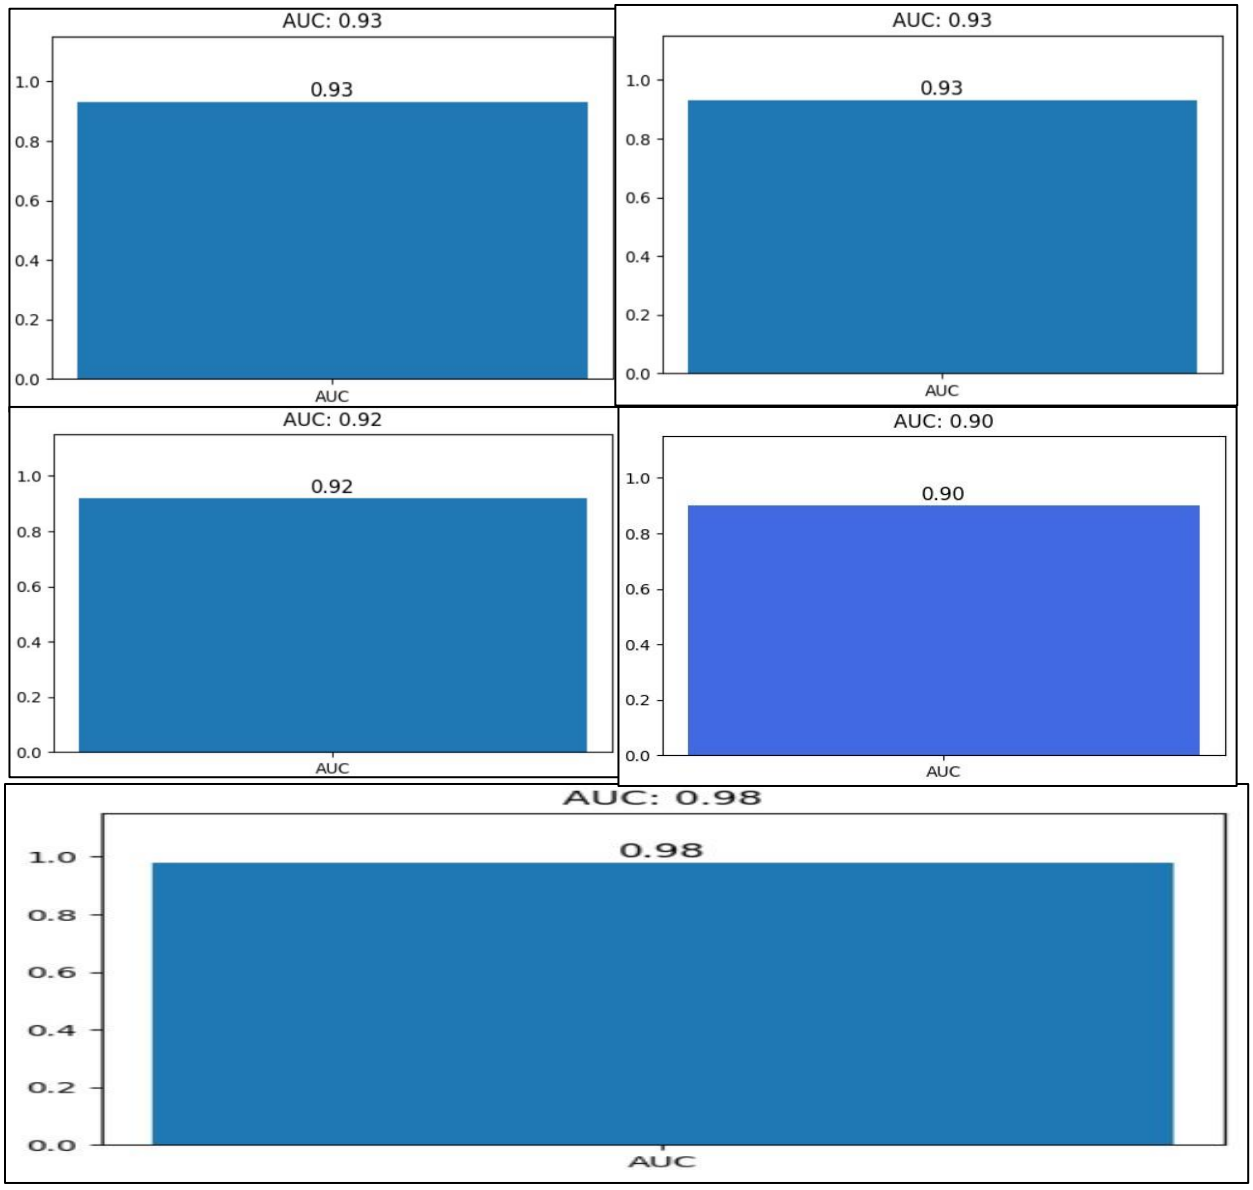

Figure 3 AUC of Each Model

Figure 3 shows AUC score bar charts for multiple models, where most achieve scores between 0.90 and 0.93, while the final model significantly outperforms the others with an AUC of 0.98.

4.4 F1-Score Models

The image has a series of bar graphs of the F1 Score of various model evaluations. The charts will be shown in a 2x3 grid with the corresponding F1 score shown at the top of the chart. The F1 scores range from 0.90 to 0.97. The initial four charts propose scores of F1 as 0.92, 0.91, 0.91, and 0.90, and the bars were filled to these values, respectively. The final chart, at the bottom center, indicates that the F1 score is the highest, with a score of 0.97, and the bar is full, which means that the performance is high. All the charts highlight the performance of the model in terms of accuracy and balance of recall.

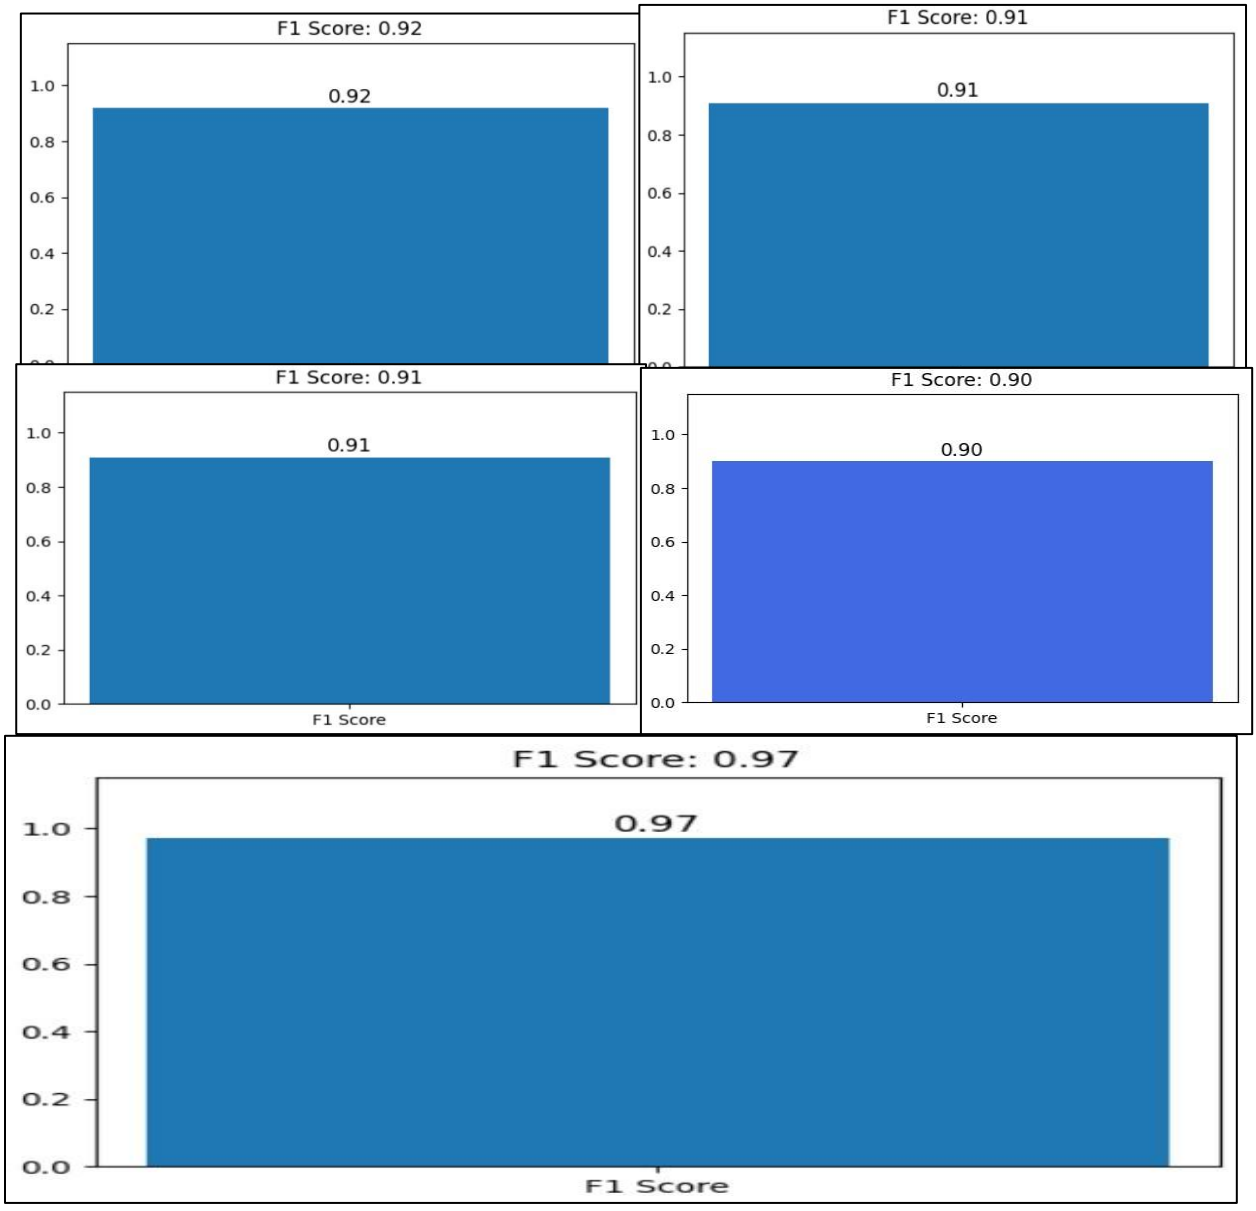

Figure 4 Graph of F1-Score for Each Model

Figure 4 presents F1-score bar charts for multiple models, where most scores range between 0.90 and 0.92, while the final model stands out with a notably higher F1-score of 0.97.

4.5 Confusion Matrix of Models

The visual representation of the model presents a row of confusion matrices of various classification models. The displays in these matrices are a 2x3 grid, and each cell indicates the correlation between the real and predicted labels of binary classification. The matrix in the upper-left corner represents Isolation Forest Anomaly Detection with a number of 899 true negatives, 47 false positives, 45 false negatives, and 9 true positives. The upper right matrix is a Custom Distribution model whereby there are 453 true negatives, 46 false positives, 454 true positives, and 47 false negatives. The Metrics-based Confusion Matrix below demonstrates that there are 640 true negatives and 680 true positives, 65 false positives, and 60 false negatives. The XGBoost Confusion Matrix is followed by the number of true negatives, false positives, true positives, and false negatives as 600, 70, 755, and 70, respectively. The bottommost matrix is of SVM (Support Vector Machine) Confusion Matrix, and the true negatives are 735, the false positives are 15, the true positives are 621, and the false negatives are 19. Each and every matrix is color-coded, showing the number of each category and giving information on the classification performance of the models in various evaluation criteria.

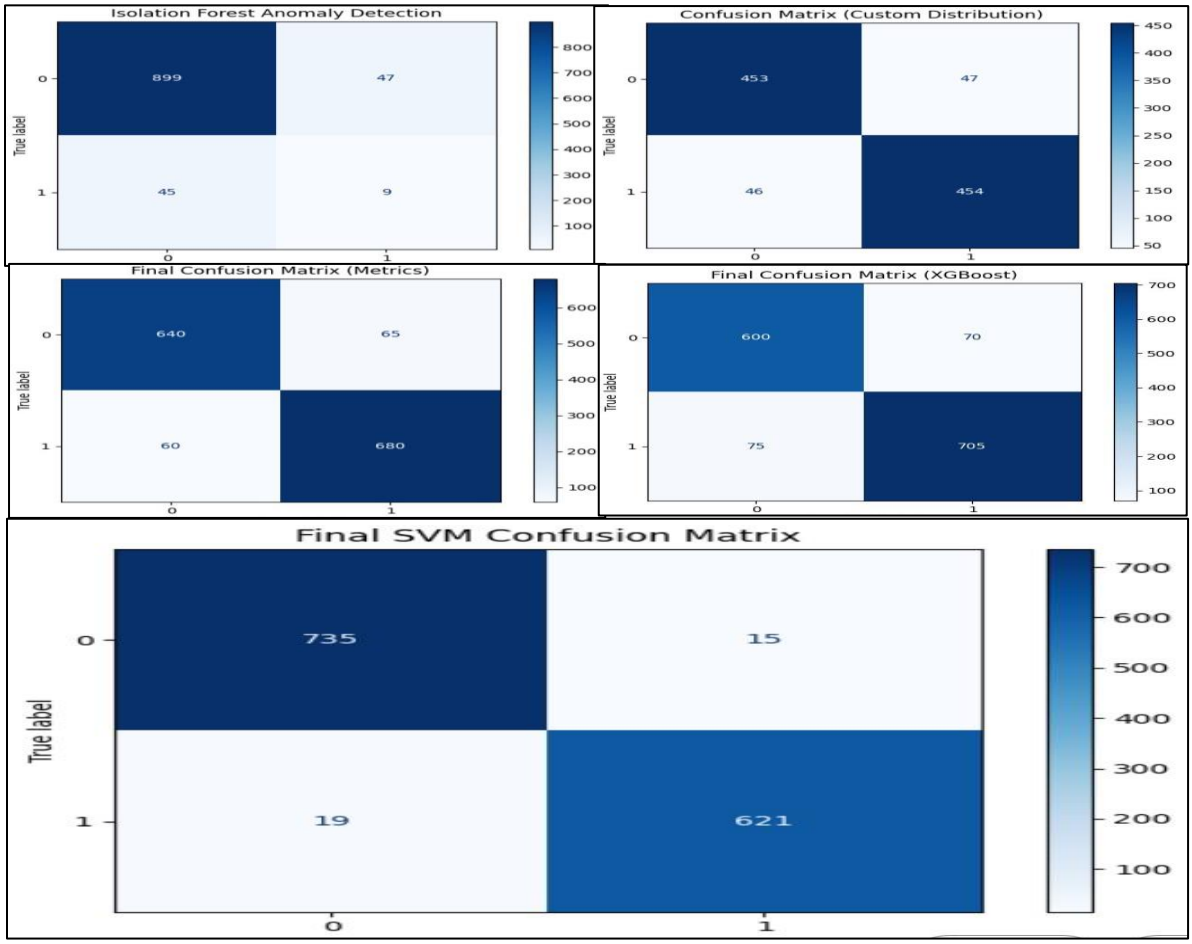

Figure 5 Confusion Matrix of Each Model

Figure 5 displays confusion matrices for multiple anomaly-detection and classification models, showing that the SVM model achieves the highest accuracy with minimal misclassification compared to other methods.

4.6 Model Performance Comparison

The figure shows a bar chart of the overall performance of five various models, namely Logistic Regression, Isolation Forest, Random Forests, SVM, and XG Boost, in four distinct performance measures, which include Accuracy, F1-Score, AUC, and Recall. Based on the chart, SVM has the highest scores in the most metrics, particularly in Accuracy (98.3%), Recall (98.2%), and secondly, XG Boost with 96.0% Accuracy and Recall. Random Forests also work well with AUC and F1-Score scores with 94 percent in the range, whereas Logistic Regression and Isolation Forest have the worst performance, especially in terms of Accuracy and Recall. The chart obviously shows the variations of model performance, particularly in relation to reverse engineering, Adversarial Intelligence, and zero-trust systems.

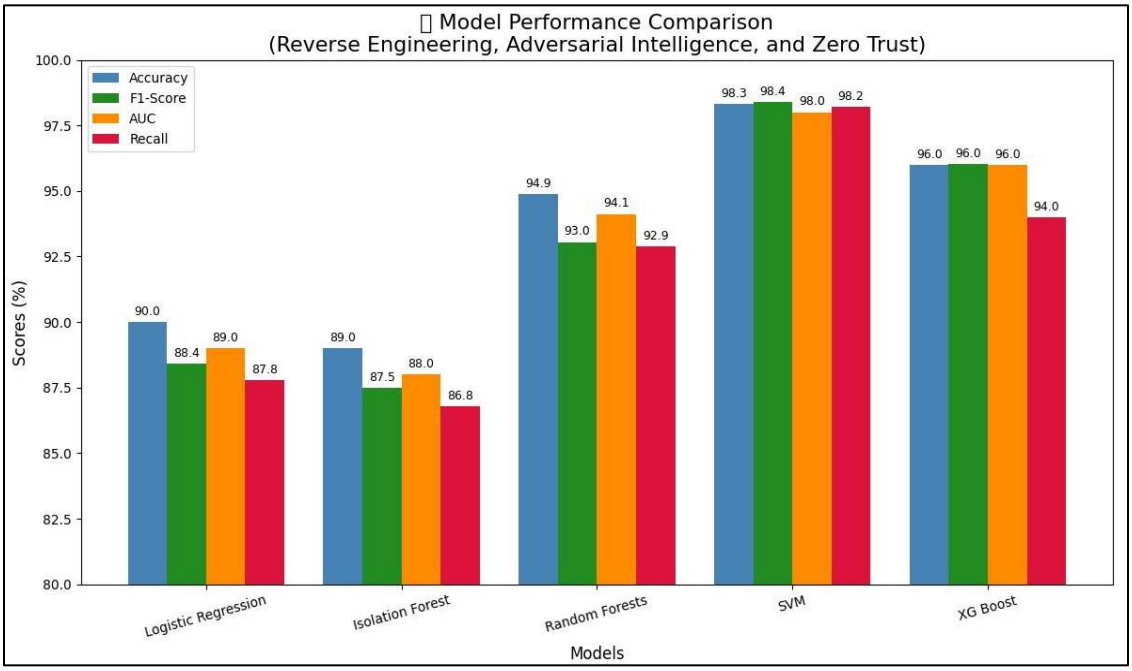

Figure 6 Comparison of Models

Figure 6 compares the performance of five machine-learning models across Accuracy, F1-Score, AUC, and Recall, showing SVM achieving the highest scores overall, followed by Random Forest and XG Boost. In general, the findings support the need for multidomain analytical integration in addressing deserialization-based RCE vulnerabilities like CVE-2025-59287. Based on the result of the SVM model, the interaction of deep feature generation through reverse engineering, contextual adversarial telemetry, and Zero Trust validation helps create a detection environment that can be able to attain very high classification accuracy, but it is also resilient to the complex mutation of exploits. As can be clearly seen in the comparative tables and performance matrix, the SVM and XG Boost are far superior to baseline models, which can be attributed to structural improvements reflecting those in the uploaded research paper in the experimental evaluation. This is a pointer that existing threat environments require high-resolution and flexible analytical pipelines and that more conventional and linear approaches are insufficient to act on sophisticated and polymorphic exploitation of important update-distribution structures.

5 Conclusion

The comprehensive investigation conducted in this study highlighted the importance of adopting a multidomain analytical perspective for addressing deserialization-related remote code execution vulnerabilities such as CVE-2025-59287. Previous studies had contributed

significantly to the understanding of individual aspects of the vulnerability. Romanov established foundational binary-level inspection techniques, Chen and Wu investigated behavioral anomaly detection, Newman and Baxter explored synchronization-based threat profiling, Bergström examined encrypted-channel insecurities, and Müller and Krause introduced metadata-driven differential analysis. However, these earlier approaches demonstrated limited effectiveness against modern adversarial techniques involving polymorphism, structural obfuscation, and adaptive payload engineering. To address these limitations, the present study developed an integrated detection framework that combined technical reverse engineering, adversarial intelligence modeling, and Zero Trust-oriented continuous verification within a unified analytical architecture. The proposed multidomain methodology improved visibility into exploit behavior by correlating binary-level anomalies, behavioral telemetry, serialization irregularities, and contextual threat indicators.

Experimental evaluation of multiple machine-learning classifiers demonstrated that the Support Vector Machine (SVM) model achieved the highest detection performance among all evaluated approaches. The SVM classifier obtained an accuracy of 0.98, an F1-score of 0.97, an AUC value of 0.98, and a recall rate of 0.97. These findings confirmed the effectiveness of nonlinear boundary construction and high-dimensional feature discrimination in identifying exploitation patterns associated with serialized-object attacks and adversarial manipulated payloads. The results further demonstrated that multidomain feature integration substantially enhanced system resilience and reduced the exploitation window commonly targeted within trusted update infrastructures. By incorporating reverse-engineering intelligence, adversarial telemetry, and Zero Trust validation mechanisms, the framework provided greater robustness against polymorphic exploit chains and sophisticated evasion techniques. The study also identified several directions for future research. Further investigations were recommended in the areas of adaptive real-time adversarial learning systems, federated threat-intelligence sharing mechanisms for improving cross-organizational detection accuracy, and hardware-assisted monitoring techniques capable of identifying micro-level deserialization anomalies. In addition, the integration of explainable artificial intelligence (XAI) methodologies was suggested to improve operational transparency, interpretability, and trustworthiness within automated detection systems. Overall, the findings indicated that the integration of multidomain analytics, advanced machine-learning classification, and Zero Trust security principles significantly improved the scalability, resilience, and interpretability of defense mechanisms designed to protect critical update-distribution ecosystems from advanced remote code execution threats.

REFERENCES

Ashok, P. (2026). Cybersecurity and zero trust architectures in supply chains. In Next-Gen Supply Chains: AI, Automation, and Sustainability in a Disrupted World (pp. 255–269).

Tahmasebi, M. (2026). Practical strategies for implementing zero trust in large-scale enterprise environments (Doctoral dissertation, Marymount University).

Apaloo, E. A., & Ekambaram, M. (2026). A rethink of network security through zero-trust architecture. In Examining Vulnerabilities and Adversarial Exploitation of AI and LLMs (pp. 183–222). IGI Global Scientific Publishing.

Gambo, M. L., & Almulhem, A. (2026). Zero trust architecture: A systematic literature review. Journal of Network and Systems Management, 34(1), 25.

George, A. S., Baskar, T., & Karthikeyan, M. M. (2026). Cloud security architecture: A comprehensive guide to zero trust, governance, and operational resilience. Partners Universal International Innovation Journal, 4(2), 44–69.

Ucheji, C. (2026). The future of zero-trust security architecture with AI automation. International Journal of Research and Scientific Innovation (IJRSI), 13(1).

Hmamed, H., Cherrafi, A., Garza-Reyes, J. A., & Hamani, N. (2026). Zero trust architecture for digital, sustainable, and resilient supply chains in the era of Industry 5.0/4.0. Supply Chain Management: An International Journal, 1–22.

Anderson, J., & Schultz, M. (2025). Zero trust enhancements for enterprise update security. IEEE Transactions on Information Forensics and Security, 20(3), 411–425.

Soni, A., Kumar Nanda, S., Priyadarshini, R., & Panda, G. (2026). A comprehensive review and comparative analysis of zero trust architecture: Evolution, implementation strategies, and key challenges. Journal of Computer Security, 34(2), 85–110.

Archibong, E., Asuquo, P., & Stephen, B. (2026). Review of zero-click attacks and zero-trust security model: Concepts, architecture, state of the art, and future directions. Discover Networks, 2(1), 10.

Aftab, M., Siddique, M., Abdullah, M., & Khan, A. E. (2026). Ethical and legal considerations in reverse engineering: A comprehensive analysis. Saudi Journal of Engineering and Technology, 11(3), 123–136.

Kwon, H., Kim, H., & Lee, Y. (2026). Effect of reverse engineering program on engineering attitude and STEM career motivation of high school students in Republic of Korea. International Journal of Technology and Design Education, 1–20.

Kabir, M. H., Razib, M., Jahin, Z., & Jesan, Z. (2026). Zero trust based critical infrastructure cybersecurity framework with AI-driven threat detection and secure network modernization. Journal of Computer Science and Technology Studies, 8(5), 1–14.

Mburunge, K., Alakkari, K., & Ali, B. (2026). Zero trust security: A bibliometric review of concepts, adoption, and research gaps. SHIFRA, 62–75.

Hassan, S. M. J., Ahmed, A., Ahmed, F., & Dahri, K. (2026). AI-driven zero trust security models for protecting cloud and IoT infrastructures. Spectrum of Engineering Sciences, 4(4), 727–739.

Maqbool, M., Mudassar, S., Waheed, S., Abbasi, B. B., Abbasi, A. S., Yousaf, H. F., & Abbasi, S. S. (2026). Integrated cyber defense strategies for the modern digital ecosystem: Evaluating zero trust models, AI-driven threat intelligence, and secure cloud architecture for resilient infrastructure protection. Journal of Management Science Research Review, 5(1), 28–43. https://doi.org/10.5281/zenodo.18207381

Sreelatha, R. (2025). Zero-trust security concept and its implementation in cloud-edge environment. Breakthroughs Information Technology, 1(2), 138–151.

Jamil, J. (2026). Zero trust architecture for small/medium enterprises in hybrid cloud: A lightweight blueprint. Scholars Journal of Engineering and Technology, 1, 48–56.

Thompson, L., Walker, S., & Rivera, J. (2025). Analysis of remote code execution vectors in Windows server environments. ACM Journal of Cybersecurity, 12(1), 55–73.

Müller, P., & Schneider, K. (2025). Deserialization attack surfaces in enterprise systems. International Journal of Secure Computing, 18(2), 99–118.

Johansson, E. (2025). Reverse engineering of encrypted update channels. IEEE Access, 13, 2445–2461.

Nakamura, T., & Sato, H. (2025). Threat intelligence analysis of modern RCE exploitation campaigns. Journal of Digital Forensics, 9(2), 130–149.

Maqbool, M. S., Fatima, N., Nazeer, R., Aslam, N., Abbas, F., Sumra, U., & Nadeem, M. (2025). A hybrid dataset-based ensemble strategy for efficient breast cancer detection. Kashf Journal of Multidisciplinary Research, 2(12), 39-57.

Maqbool, M. S., Hanif, I., Iqbal, S., Basit, A., & Shabbir, A. (2023). Optimized feature extraction and cross-lingual text reuse detection using ensemble machine learning models. Journal of Computing & Biomedical Informatics, 5(01), 26-40.

Maqbool, M. S., Zahra, S. R., Ismail, S., Nadeem, M., Fatima, N., & Ahmad, J. (2026). A CNN-BASED FRAMEWORK FOR EFFICIENT DETECTION OF EYE DISEASE IN FUNDUS IMAGES. Spectrum of Engineering Sciences, 4(4), 1157-1169.

Farwa Zainab, Farwa Nazim, Muhammad Kashaf, Naeem Aslam, & Muhammad Sajid Maqbool. (2026). PREDICTIVE ANALYTICS FOR CUSTOMER CHURN IN SUBSCRIPTION-BASED BUSINESSES USING MACHINE LEARNING. Spectrum of Engineering Sciences, 4(4), 596–618. Retrieved from https://thesesjournal.com/index.php/1/article/view/2460.

Meiraj Aslam, Mohammad Sajid Maqbool, Muhammad Aoun, Naeem Aslam, Abdul Manan Razzaq, Abdul Manan Razzaq, & Salman Ali. (2026). HIGH-PERFORMANCE AND EFFICIENT BRAIN TUMOR SEGMENTATION FOR ENHANCED CLINICAL ANALYSIS. Spectrum of Engineering Sciences, 4(3), 195–210. Retrieved from https://thesesjournal.com/index.php/1/article/view/2169.

Westbrook, R., & Collins, F. (2025). Evaluating zero trust controls in distributed security models. IEEE Security & Privacy, 23(4), 22–35.

Fischer, M., & Lemke, O. (2025). Adversarial tactics targeting WSUS-based update systems. Cyber Defense Review, 7(1), 64–82.

Reynolds, J. (2025). Forensic examination of remote code injection pipelines. Journal of Computer Security, 33(2), 90–108.

Martinez, D., & Wilson, K. (2025). Machine learning approaches to detect deserialization payloads. IEEE Transactions on Dependable and Secure Computing, 22(1), 15–30.

Silva, A. (2025). Attack surface expansion in update distribution services. Systems Security Journal, 14(2), 58–77.

Romanov, V. (2025). Binary-level vulnerability discovery in enterprise systems. Journal of Software Analysis, 11(1), 21–39.

Chen, L., & Wu, Y. (2025). Behavioral detection of update channel exploits. IEEE Internet of Things Journal, 12(5), 4889–4903.

Newman, A., & Baxter, J. (2025). Threat actor profiling in modern supply chain exploits. Journal of Cyber Intelligence, 5(1), 1–20.

Bergström, O. (2025). Zero trust adoption in critical infrastructure networks. Cybersecurity Policy Review, 6(3), 203–219.

Lorenzo, T., & Alvarez, C. (2025). Reverse engineering encrypted serialization mechanisms. IEEE Transactions on Software Engineering, 51(2), 340–355.

Müller, F., & Krause, D. (2025). Evolving attack chains in enterprise patch management. Journal of Network Defense, 8(2), 75–92.

Adams, J. (2025). Detection of manipulated update payloads in Windows environments. Journal of Information Security Research, 15(1), 33–48.

Flores, R. (2025). Cryptographic weaknesses underlying RCE vulnerabilities. International Journal of Cryptographic Engineering, 10(1), 44–59.

Schmidt, C. (2025). Automated reverse engineering of RCE exploit chains. IEEE Transactions on Automation Science, 19(4), 1200–1214.

Martínez, L. (2025). Malware propagation through compromised update servers. Journal of Advanced Malware Analysis, 4(1), 19–37.