AI-Enhanced DevSecOps for Secure and Compliant CI/CD Pipelines

Iveri Jajanidze
Georgian Technical University
jajanidze.iveri24@gtu.ge

Abstract

Continuous Integration and Continuous Deployment (CI/CD) pipelines have revolutionized software delivery by enabling rapid iterations and reducing release cycles. However, this acceleration of deployment also magnifies risks: insecure code can propagate more quickly, compromised dependencies can infiltrate multiple environments, and automated release processes broaden the attack surface. Traditional, rule-based defenses struggle in such dynamic environments.

This paper introduces an AI-Enhanced DevSecOps framework that strengthens pipeline security through the integration of Explainable AI (XAI), Federated Learning (FL), adversarially resilient malware detection, and AI-augmented defenses against social engineering. By combining these techniques with DevSecOps principles, organizations can achieve both resilience and compliance with regulations such as the GDPR, NIST SP 800-53, and the EU's Digital Operational Resilience Act (DORA). Real-world case studies, including SolarWinds, Mirai Botnet, and phishing simulation research, demonstrate the necessity and practical benefits of such integration. Results highlight measurable improvements in detection accuracy, mean time to resolution, and overall trust in automated security decisions.

Keywords: DevSecOps, CI/CD Security, Explainable AI, Federated Learning, Adversarial Attacks, Compliance

1. Introduction.

In modern software engineering, CI/CD pipelines have become indispensable for delivering code efficiently. Continuous Integration allows developers to merge changes frequently into shared repositories, while Continuous Deployment automates the process of pushing those changes into production environments. This shift has fostered agility, but it has also exposed development ecosystems to unprecedented risks.

The SolarWinds compromise of 2020 exemplifies how software supply chains can be exploited, with attackers embedding malicious code into a vendor's update mechanism, thereby affecting thousands of organizations [1]. The Mirai Botnet of 2016 demonstrated how seemingly trivial vulnerabilities, such as default credentials, can be weaponized at scale [2]. These incidents illustrate that static, pre-defined defenses are insufficient in a landscape where adversaries adapt rapidly and exploit automation itself as a vector of attack.

DevSecOps emerged as a strategy to embed security into every phase of the CI/CD pipeline. Yet this approach must now confront adversaries who themselves employ artificial intelligence to craft polymorphic malware, evade detection, and automate reconnaissance. The logical evolution is AI-Enhanced DevSecOps, a framework that integrates adaptive, explainable, and privacy-preserving AI methods directly into the pipeline.

2. Methodology.

The AI-Enhanced DevSecOps framework builds upon traditional DevSecOps principles but introduces intelligent automation at key stages of the pipeline. The methodology rests on four interrelated pillars: transparency, collaboration, resilience, and human-awareness integration.

Explainable AI (XAI): Machine learning models used in anomaly detection or risk prediction are often criticized as 'black boxes.' Explainable AI addresses this by making their reasoning interpretable. Techniques such as LIME and SHAP allow developers and auditors to understand why a build or deployment was flagged, thus ensuring accountability in accordance with regulations like the GDPR's 'right to explanation' [3].

Federated Learning (FL): Security data is often distributed across different departments or even organizations, yet pooling this data can create privacy risks. Federated Learning allows the development of shared models without exposing raw data. Instead, only aggregated updates are exchanged, which preserves confidentiality while enabling collaboration on threat intelligence [4].

Adversarial Resilience: AI systems themselves are vulnerable to evasion, poisoning, and adversarial examples. To counter this, the framework incorporates adversarial training and ensemble learning [5].

AI-Augmented Social Engineering Defense: Phishing campaigns, enhanced by AI, now generate messages that are contextually tailored and linguistically convincing. Combining natural language processing with phishing simulation exercises, the framework reinforces resilience by both detecting malicious communications and improving user awareness [6].

3. Discussion.

Integrating AI into CI/CD security offers clear benefits but also raises important considerations. On the positive side, automation accelerates the detection of anomalies in logs, APIs, and builds, reducing the mean time to resolution by almost half [7]. Furthermore, XAI provides interpretability that not only assists developers in debugging but also satisfies the stringent requirements of auditors and regulators.

However, the introduction of AI does not eliminate challenges. Adversarial machine learning represents a growing field of attack where threat actors deliberately exploit weaknesses in models [5]. Addressing these vulnerabilities requires constant retraining, monitoring of model drift, and a governance framework that balances innovation with accountability [2].

Another challenge lies in resource constraints. Federated learning and ensemble methods demand significant computational capacity, which may limit adoption in smaller organizations. This creates a gap between enterprises that can operationalize AI-based DevSecOps at scale and those still reliant on traditional security measures.

4. Case Studies.

The practicality of AI-Enhanced DevSecOps becomes evident when applied to high-profile incidents.

The SolarWinds attack (2020) was characterized by insertion of malicious DLL files into build artifacts. Anomaly detection models trained to analyze code provenance and build integrity could have identified unusual insertions before distribution [1].

The Mirai Botnet (2016) exploited weak IoT credentials to mobilize a botnet capable of massive DDoS attacks. AI-driven anomaly detection applied to login behavior and network traffic could have signaled the abnormal surge in device activity and enabled earlier intervention [2].

Recent phishing simulation studies (2024) have shown that simulated training significantly decreases user vulnerability to social engineering [6]. When coupled with AI-based natural language analysis of emails and messages, organizations can achieve a dual line of defense-both technical detection and improved human response.

5. Results.

Empirical research reinforces the benefits of AI-driven DevSecOps. Jajanidze [1] reports that anomaly detection models such as autoencoders and Isolation Forests achieved AUC scores above 0.90 in detecting pipeline anomalies. Mean time to resolution was reduced by approximately 45 percent compared to traditional systems, highlighting tangible operational gains.

Explainable outputs enabled clearer audit trails, supporting compliance with GDPR and DORA. In parallel, phishing simulation research demonstrated a reduction in susceptibility rates of more than 30 percent when users were trained with AI-supported tools [6].

6. Conclusion.

The convergence of AI and DevSecOps represents a decisive step toward securing modern software delivery. CI/CD pipelines, while indispensable for agility, expose organizations to risks that static security controls cannot mitigate. AI-enhanced techniques-explainable models, federated learning, adversarial defenses, and social engineering countermeasures-offer a comprehensive response.

The evidence from case studies and empirical evaluations shows that such approaches improve detection accuracy, accelerate response times, and reinforce compliance. At the same time, challenges remain: adversarial threats evolve continuously, and resource constraints may hinder adoption in smaller organizations.

References:

[1] I. Jajanidze, The Use of Artificial Intelligence in CI/CD Systems, Georgian Scientific Journals: Computer Science & Telecommunications, 2025.

[2] C. Cath, et al., Artificial Intelligence and the ‘Good Society’: The US, EU, and UK Approach, Science and Engineering Ethics, vol. 24, no. 2, pp. 505–528, 2018.

[3] D. Gunning, et al., XAI-Explainable Artificial Intelligence, Science Robotics, vol. 4, no. 37, 2019.

[4] H. Yang, et al., Federated Machine Learning: Concept and Applications, ACM Transactions on Intelligent Systems and Technology, vol. 10, no. 2, pp. 1–19, 2019.

[5] D. Ucci, L. Aniello, and R. Baldoni, Survey of Machine Learning Techniques for Malware Analysis, Computers & Security, vol. 81, pp. 123–147, 2019.

[6] A. Bichnigauri, I. Kartvelishvili, O. Shonia, D. Bichnigauri, and O. Gudadze, Strengthening Cyber Defenses – The Crucial Role of Phishing Simulation in Modern Security Strategies, Defence and Science, no. 3, 2024.

[7] I. Kartvelishvili, G. Kuchava, Optimization of Software Delivery in DevOps with CI/CD, Proceedings of the International Scientific-Practical Conference, Georgian Technical University, 2024.