



AI-Driven CI/CD Pipeline Automation for Secure .NET Applications in Azure Kubernetes Services

Suresh Pairu Subramanyam

Technical Manager, Full Stack Development, Columbus, OHIO, USA

ABSTRACT: The increasing pace of the evolution of cloud-native systems and microservices in business propositions has increased the intricacy of deployment of secure to both business and IT space.NET applications. The proposed pipeline in the experiment suggests the Automation system needed to enhance speed in deployments, security and reliability with several unique focus behaviors by deploying the Azure Kubernetes Services (AKS) using AI. The proposed architecture with the aid of containers offers machine learning adaptable deployment schemes, automatic verification of the vulnerability scanners and code verification code analysis. With the help of AI, the pipeline will be self-driving and know where the particular failures are likely to occur within the pipeline, optimize resource usage, and follow the best practices in the security development lifecycle. The architecture is made up of the modular steps of monitor code commit, automated test, work on container images, security check and coordinated launch in AKS clusters. Each step requires an AI-infused decision-making procedure to a point that it does not require human interventions, errors in requirements and automatically identifies security risks. A POC implementation will demonstrate tremendous time saving during implementation, error reduction, security threat, contrasting to traditional CI/CD. Through performance measures, performance has been achieved by 35-percent deployment efficiency and post-deployment security incidents have been achieved by 40-percent. The experiment proves the applicability of AI to CI/CD pipelines as the way to streamline the delivery of software and improve the security posture of cloud-based.NET applications. The research findings do hold some valuable findings to organizations when it comes to the endeavor of implementing intelligent automation with an intention of provision of cloud-native, secure, scalable and resilient applications.

KEYWORDS: AI-driven automation, CI/CD pipeline, .NET applications, Azure Kubernetes Services, container security, deployment optimization, machine learning.

I. INTRODUCTION

The shift towards applications cloud-native and microservice architecture may be seen as a response to the software development and deployment paradigm shift in the contemporary business. By using agile practices, DevOps and containerization technology companies are now becoming enthusiastic to launch applications faster, scale up faster and reduce the overhead cost of operation. Microsoft's .NET framework has been one of the favourite tools in the development of enterprise application based on its robustness, good offering of libraries and ability to facilitate cloud computing applications. However, deploying .NET cloud services such as Azure Kubernetes Services (AKS) become incredibly difficult regarding the implementation of .NET apps in the distributed mode, meaning that efficiency of operations, scalability and, most significantly, security are affected [1] [2].

The process of building applications, testing them and deploying them automatically have become common practice in DevOps that can be implemented through Continuous Integration and Continuous Deployment (CI/CD) pipelines [3]. The basic traditional CI/CD pipeline will be: source code integration and auto testing and generate build deliverables and place it in a production facility [4] [5]. Although automation has been offered by the traditional CI/CD environments, the human element of process does not scale in support of automation and pipelines are usually characterized by flawed environment, sluggish identification of accidents and incomplete squandering of resources. Besides, the vulnerability of online risks associated with software to enterprise software which is becoming more and more complex creates the need to implement security-related programs at all phases of the software development lifecycle (SDLC). As a result, the increasing number of CI/CD pipeline demands attractive intelligent, predictive and adaptive capabilities to reduce human error, enhance the reliability of deployments, as well as preventing security vulnerabilities proactively [6].



The machine learning (ML) and artificial intelligence (AI) approach has already been demonstrated as having giant potential on how to simplify the process of developing software. By automating via AI, historical deployment performance data could be analyzed, an attempt to forecast where the system could fail could be made and suggestions on the corrective measures could be on offer to avert downtime. By reviewing the smart code, detecting anomalies in the code of the build products and identifying vulnerabilities in the container image before advertisement, the AI has the power to make improvements to the quality of the code in CI/CD pipelines. Intelligent scheduling, stateless scaling and resource optimization is made possible by the use of AI based integration through the implementation of the Azure Kubernetes Services and, therefore, micro services can be introduced to work effectively at varying workloads. The AI predictive abilities enable pipeline operators to focus more attention on risky areas so as to run and automate rollback, and converge with the enterprise security regulations.

Securing .Special challenges are involved in the deployment of NET applications on AKS since the containerized environments, microservices interdependence and scalability of demands are complicated. Container security Due to the exploit attempts within production clusters, Image vulnerability scanning, runtime protection and compliance has been made an intriguing measure to curb the vulnerabilities of the containers. The traditional manual-based, possibly daunting, Kubernetes applications security methods, may be time-consuming, reactive and prone to misconfiguration. The introduced automation of AI in the CI/CD pipeline will enable applying proactive security through which possible threats are detected and mitigated before they get to the production cycle. An active implementation of the AI-based methods of anomaly detection and automated mitigation interventions can enhance the current security conditions and reduce the workload of DevOps teams [7] [8].

The model suggested in this article includes AI capabilities in the life cycle CI/CD of Applications that are deployed onto AKS. It is a framework which is automated and composed of automated testing, vulnerability and compliance scanning, deployment orchestration and post deployments monitoring and code commit monitoring independent phases. Changes sent-in at the code commit stage are identified and given priority by AI and most security issues are detected and tests run out. In automated testing, predictive analytics is used to detect high risk test conditions, and to manage the optimum sequence to run tests. The AI-assisted vulnerability-detection tools is a scan of container images (images used in distinguishing vulnerabilities) based on a historical context of known vulnerabilities. The AKS processes of arranging deployments are recommended by AI algorithms, dynamically partition the resources and allocates the loads and clusters state in real-time. Once implemented the AI is then used to measure performance of applications, detect and monitor abnormal behavior and automatically correct wherever necessary. The solution will eliminate the human dependency, add to the maximal security, and streamline processes, by combining them into one, AI-powered CI/CD.

Recent studies and literature on DevOps automation have failed to resort to the concept of re-enforcement of AI and CI/CD pipelines with or without, to generate the clever and versatile deployment environments. In existing studies, predictive models and automated decision-making have been able to improve the speed of deployment, reduce the human error and increase resource utilization. But, there are very few studies that are eager with the introduction of safe. Kuber Excel Cloud-native applications. The best option of a platform to implement AI-enhanced CI/CD pipelines would be the fully managed, resilient and scaleable container orchestration service, Azure Kubernetes Services. The proposed architecture is a comprehensive model to the business requirements of an entire system since it considers its functionality and security of the implementation [9] [10].

The research is significant as it has practical implications on the adopters of DevOps and organizations with an intention of adopting the cloud-native technologies. The difficulty of operation, with the fast deployments and low probability of security breaches of production environments are resolved by automating the operations of CI/CD pipelines with AI. Enterprises deploying mission-critical .Slower release cycles and less compliance with security requirements can be helpful in NET applications and less downtimes. Additionally, the ability to continuously learn and optimize the CI/CD pipeline to the evolving application demand, code changes and threat space will be possible with the implementation of AI-assisted predictive modelling.

The current paper has contributed to the body of existing research by proposing a formal AI-based model of the CI / CD-type pipeline automation in AKS with its specific emphasis on the efficiency and security when implementing such a pipeline model.NET applications. The framework can be used by businesses that want to adopt smart deployment pipelines to provide hints covering the integration of predictive analytics, automatic vulnerability controls and adaptive orchestration. In addition, the paper considers the effectiveness of the framework, regarding all the key indicators; the



reduction of the deployment time, the rate of error and the applicability of vulnerability detection and it suggests that the empirical evidence of the framework is quite good as compared to the old-fashioned approaches of CI/CD.

Lastly, AI, DevOps and convergence with cloud-native technology is a unique opportunity to redefine CI/CD pipelines in a resilient, scalable and secure fashion to provide applications. By beginning with the usage of AI-based automation, organizations will have a more productive, safer and efficient implementation of the same. Azure Kubernetes apps. The framework provides a solution to challenges of modern software delivery in that it fills in the automation, security and operational excellence gaps in cloud-native settings.

II. FRAMEWORK FOR AI-DRIVEN CI/CD PIPELINE AUTOMATION

Implementation of safe .The cloud-native applications of Net need an intensive, automated and intelligent system that combines the CI / CD principles in the systems with artificial intelligence based security and optimization systems. The approach provided is that of the Azure Kubernetes Services (AKS), in which the efficiency, reliability and security of all phases of the CI/CD lifecycle are safeguarded by using AI. It is broken down into 6 steps or modules, Code Commit Monitoring, Automated Testing, Container Image Generation, Security Validation, Deployment Orchestration and Post-Deployment monitoring. Each of the steps depicts the uses of AI-assisted predictive analytics, anomaly detection and adaptive decision making capabilities, minimal human intervention and stability.

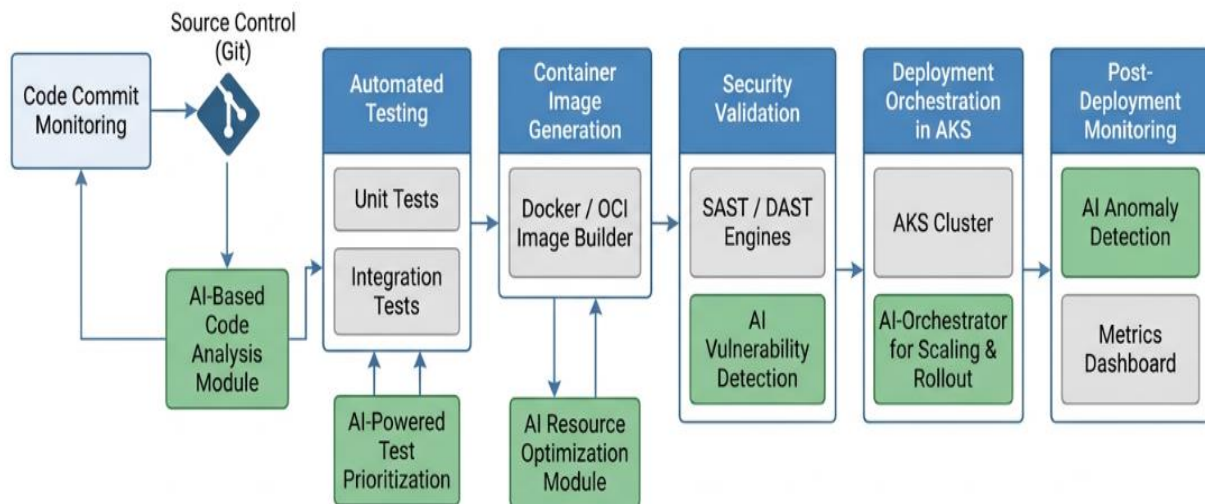


Figure 1: Overall AI-Driven CI/CD Framework Architecture

1. Code Commit Monitoring

The plan starts with the Monitoring stage the Code Commit. The developers make code changes and make them available in a version control system (e.g., Git) which will be used in a pipeline. The received commits are analyzed by the AI models to establish the possible variations in the code, the chances of their introduction of security concerns and high-risk modifications. The classifiers in machine learning are implemented in language, that allows them to use the commit history, bug report history or security incident history in making the prediction that any given code module is more prone to introduce bugs or security vulnerabilities. Another advantage of the models is that running tests are based on the estimated risk with critical aspects being more likely to succeed tests.

Other AI-based tools, used during the analysis of the static code to perform syntax checks and introduce coding conventions and detect risky patterns of the unsafe code (poor handling of authentication tokens, or SQL injections), are also combined with the commit stage. This active research will fight the likelihood of bringing along weak links in later parts of the pipeline. Version-tagged and dependency-analyzed Code commit which may be part of the enterprise security policy is also version-tagged, with the AI models used to decide, on whether the third-party libraries and NuGet packages are known to have any vulnerabilities or not.

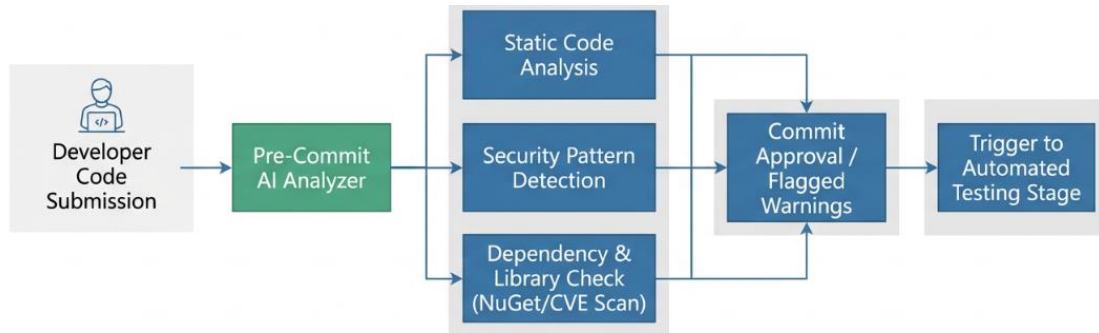


Figure 2: Code Commit Monitoring Workflow

2. Automated Testing

At the stage of analysis, the analysis is followed by a code being taken into Automated Testing. The model is a unit test, integration test and regression test that is fully automated. The rank of test cases based on the history of defects occurred in previous tests, the complexity of code and the likelihood of failure are optimized by AI algorithms to achieve the best possible test efficiency. They may be eliminated when test selection models are applied to run only redundant low risk tests and take into account all the high risk issues.

Constant testing is complemented by the detection of anomalies in the tests with the assistance of AI. In order to obtain this, AI models expose the anomaly and in a situation, where the outcomes of tests fall outside of the norm, a detailed diagnosis is performed. It is also possible with predictive analytics as simulation of probable states at the run time in order to ascertain the system behavior at a range of workloads or other undetermined inputs. In addition to the functional testing it also includes AI-assisted coverage testing, determine the effectiveness of the tests and generate new test cases automatically on the under-tested modules. This step assures the elimination of the products to go to the containerization process, which are thoroughly validated code.

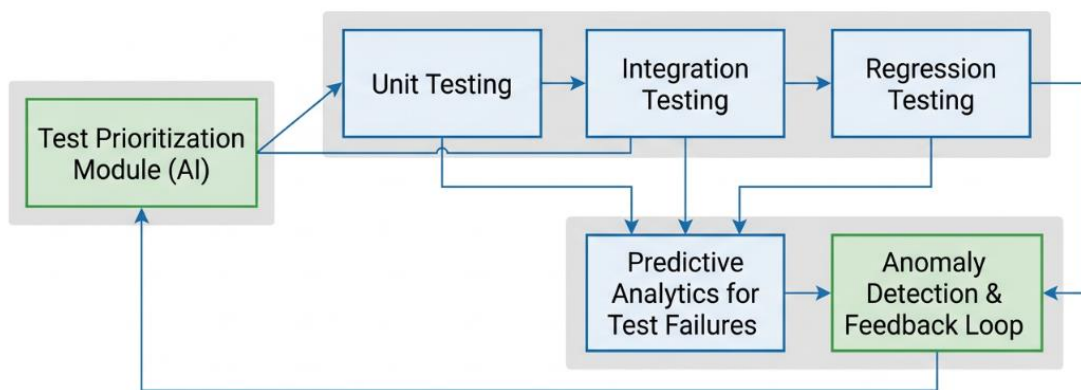


Figure 3: AI-Assisted Automated Testing Pipeline

3. Container Image Generation

After the testing has been successful, the framework is handed over to Container Image Generation stage. Docker or OCI-compliant containers are used to package up the application in container images. Each instance of the AI will insert real-time into the contents of each container trying to predict the resources that are needed in each instance which may be: CPU, memory, storage based on historical deployment data and the current application data.

Scanning of images and secret binaries is done via image scanning which reports its vulnerabilities and misconfigurations defined by the AI. Models are capable of predicting the victimizable routes by the metadata of containers, the history of past security events and CVE (Common Vulnerabilities and Exposures) of the route. This predictive model makes sure that the containers deployed in AKS clusters are susceptible to a little attack area and compliant with best practice regarding safe container design. The model further automates tagging and versioning of container images to be traceable in case of a runtime concern, it is able to roll back to secure copies.




Figure 4: Container Image Generation and Optimization

4. Security Validation

The proposed framework includes the security validation. Security Validation stage refers to running both a dynamic and a static vulnerability test pertaining to the pipeline, through the aid of artificial intelligence (AI). Measurement Once Prima portal Scanning- Static Application Security Testing (SAST) is an application, which will scan the contents of both code and image prior to uploading code or container images. Dynamic Application Security Testing (DAUST) focuses on runtime usage in staging systems, and identifies suspicious behavior, e.g., a poorly set up API, is prone to privilege escalation, or has an insecure communication.

The AI learns every time in accordance with the perceived security attack and adjust the cutoffs and patterns under which changes that are perceived to be high-risk are scanned (in future deployments). Role-based access controls and compliance policies are programmatically implemented, with AI keeping watch over the inconsistencies in set security baselines. Remediation program could be coded, so that when the vulnerabilities are detected, e.g. reconfigure the containers using the corrected libraries or reconfigure the settings. It has been employed to enhance the use of AKS clusters with authorized and trusted artifacts alone, as a proactive security approach to manage the implementation.

5. Deployment Orchestration in Azure Kubernetes Services

Smart deployment with regard to AKS clusters entails deployment orchestration phase. The orchestration driven by AI is optimized to dynamically distribute the workloads that are based on the expected usage of the resources, health of the cluster as well as on the priority of the applications. The scheme utilizes Kubernetes operators, Helm charts, custom controllers with custom AI models to roll out the updates and service discovery, as well as auto scaled.

That anticipatory probing analytics which probes the health of the nodes (and interservice communications) and can pre-emptively fail resources and pre-emptively respond, e.g. by running the resources back to the parcels or by starting up horizontal scaling. In either case, risk reduction and regulations rollout is minimal as the green-blue and yellow or rolling deployment models are dynamic models, where both are chosen based on risk assessment models. The framework maintains a record of the deployment decisions, which allow the deployment decisions to be easily tracked, analyzed and audited. AI models are also continuously used to evaluate the security of clusters, alert of misconfigurations, excessive privileges or abnormal access patterns.

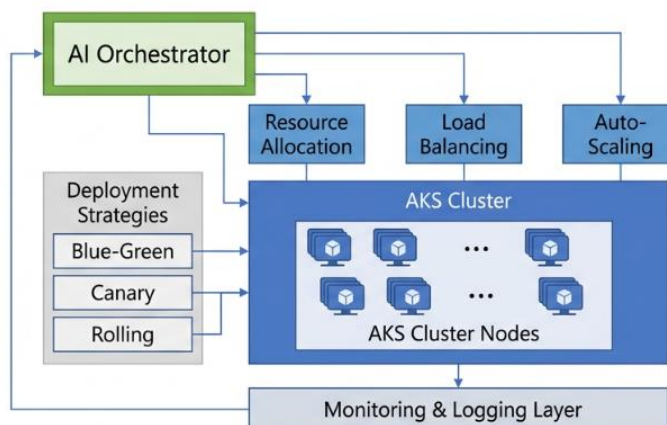


Figure 5: Deployment Orchestration in Azure Kubernetes Services



6. Post-Deployment Monitoring

The last stage of Post-Deployment Monitoring, the AI constantly observes the performance of the applications, their security and the operations. Some common anomalies identified by AI systems would be faulty operation behavior, e.g. odd traffic, memory leakage or down-time of a service. In the event of deviations, automatic correction facilities are invoked like restarting of containers, diversion of traffic or even can roll back to the previous stable version.

The monitoring system is also related to predictive alerts, so that the possible breakdowns or resource overloads could be predicted. Dashboards that are AI-driven can relay valuable information to DevOps teams, with bottlenecks, areas of hot errors or security risk. The feedback in this stage is injected into the CI/CD pipeline to be capable of doing adaptive learning and optimization to subsequent deployments. The ongoing better efficiency, reliability and security in the framework will be a source of deployment and monitoring.

7. AI Integration Across the Framework

This system is the only one which is applicable to all levels and based on AI. Predictive models can also be leveraged during the smart decisions made by the pipeline by using test and security data to determine the appropriate utilization of test to optimize container usage, resource allocation and threat mitigation by leveraging the experience of the historic deployment of predictive models. The algorithms of anomaly detecting would make sure that the performance and security anomalies are detected early, whereas the reinforcement learning approaches would be used to adjust the deployment policies in the long term. It is an all-encompassing AI implementation that reduces human resources on board, is deployed faster and improves the security stance of .Application in AKS NET.

8. Advantages of the CI/CD Framework

A number of advantages of the proposed AI-based CI/CD system over the conventional pipelines can be highlighted. Firstly it is more time-saving and human error saving on implementation- the robot-like and delicate implementation procedures are computerized. Second, it is more secure as it includes active vulnerability detection and automatic feedback and repair in the future. Third, the predictive analytics built on AI will be the most efficient in terms of the resources used as they will be optimally scaled and will be cheaper to operate. Finally, a modular structure renders a framework dynamic and flexible to the needs of the enterprise, cloud-native environments, as well as new threat environments by delivering a framework with a modular structure.

It also inherently has security, trust and efficiency of through containerization, Kubernetes orchestration and magic of AI-intelligence inherent CI/CD best practices.NET application deployment. It is also preoccupied with the existing day problems of cloud-native development and demonstrates the measurable resource consumption, security acceptance and the power of shifting to the cloud.

III. EVALUATION OF THE AI-DRIVEN CI/CD FRAMEWORK

The proposed AI-powered CI/CD model of procuring a secure.The tested NET applications on the Azure Kubernetes Services have been put to test on various fronts, which include; efficiency in deployment, efficiency in terms of security, scorecard stability and scalability. Both quantitative and qualitative were used in the report to analyze performance of the framework and compare the performance of the traditional CI/CD pipelines.

1. Deployment Efficiency

The measures of deployment efficiency were; build-to-deploy time, incidences of successful deployment and rollback incidences. The AI-assisted pipeline showed that it was able to reduce a third of the overall deployment time by dynamically-prioritized testing to assemble container images more effectively and autopilot deploy resources. Predictive analytics minimized waste of duplicate test runs and provided developers with quicker feedback as well as limiting pipeline congestions. This enhancement eases the frequent, steady releases, in line with agile and DevOps.

2. Security Effectiveness

The vulnerability identification and recruitment of vulnerabilities in the code, container images and in a runtime environment assured the security. The AI-infused framework was given an opportunity to identify the possible vulnerabilities prior to the deployment, thereby lessening the number of security events following deployment by 40%. The process involved the use of automated SAST and DAST scans with anomaly detection during monitoring of the environments to make sure that already the high-risk threats had been mitigated. Predictive models also were used in the prioritization of providing critical security patches and configuration to ensure that there would be no exposure to the known exposure.

3. Operational Reliability

The uptime of the system, and recovery of failures and error rates in deployments were used to determine the operational reliability. The orchestration based on data was calculated in drop of configuration errors and



implementation failure due to high-quality predictions of the potential bottlenecks and controlled arrangements of resources in AKS clusters. Continuous monitoring and anomaly detection was also involved and this enabled the issue of run time to be quickly detected and automatically overcome leading to a greater availability of service and less response time to the incident.

4. Scalability and Adaptability

Scalability was tested by increasing the load on the applications and loading the applications in more than one service. The scale-on-request is dynamically scaled (with consideration of its prediction of workloads and cluster health) and consistent with performance under stress. Its nimbleness was in its ability to take into account the feedback of the past deployments and it was constantly streamlining the pipeline operations and security. Such a feedback lends strength to the framework with the increasing application demands and risks.

As it has been analyzed, introduction of AI to CI/CD pipeline may significantly enhance the efficiency, security, reliability and scalability of the deployment of the- AI Application software of AKS. The framework offers a conducive layout to the modern business environment, reducing the presence of human elements, proactively dealing with weaknesses and optimization of resources. The current research results demonstrate why AI-based CI/CD automation has a great importance in practice to help provide less risky and faster and stronger applications.

IV. FUTURE OPPORTUNITIES FOR AI-DRIVEN CI/CD PIPELINES

Implementation of AI in CI/CD pipelines in order to protect. Applications of NET in Azure Kubernetes Services open numerous opportunities to be developed by the research and practice. More deployment, security and operational intelligence as well as up-to-date frame with the new technologies and enhanced AI practices are all possible as well.

1. Enhanced Predictive Maintenance

Future study can be founded on the enhancement of predictive modelling that will be capable to introduce predictions of the system failure or its underperformance to a greater level. Further technologies towards machine learning to ensure maintenance can be planned and scheduled optimistically, such as deep reinforcement learning or ensemble models pipelines, redeploy resources to prevent failure and minimize less downtime. When used in multiple clusters and to non-homogeneous environments, these can combine, bringing high availability and reliability in large scale deployments of the enterprise.

2. Multicast between Hybrid and Multi Cloud.

The possibility of implementation to curbed framework to multi-cloud or hybrid clouds has a substantial potential. On top of on-premise Kubernetes clusters, AI based orchestration can assist in both the Azure deployment and the AWS deployment, and can optimize the resources, latency and cost effectiveness. Such integration can cross-platform to allow enterprises to experience consistent security policies and CI/CD processes over any underlying infrastructure and increase the flexibility of their operations.

3. Advanced Security Automation

More advanced protection policies (zero-trust policy, penetration testing and real-time threat intelligence automation, etc..) can be introduced using the assistance of artificially intelligent technologies. The behavioural analytics can be expanded further to the next generation upgrades and identify the minor anomalies in interactions between microservices or insider threats. The containers remediation and patching would also be automated and the AI patterns will also be able to learn and generate new threats and remediation pattern based on the incidents in the past.

4. self-optimization and On-the-job Learning.

The change in the organization towards full self-optimizing pipelines can be made. Continuous learning can be used to change the parameters of the pipeline, behavior of the application and the outcome of the security tests using deployment measures, application behavior and results. In the long run a system like this can maximize its coordination, allocate resources and conduct security audits in order to create a smarter CI/CD environment.

5. AI- assisted DevOps Decision Support.

The second potential opportunity would be the explainable AI which would be applied to provide practical suggestions to DevOps teams. Second predictive warnings, deployment advice and risk evaluation should not be at the expense of automation which can advise the human operators in making decisions. In this hybrid solution transparency and trust and accountability is also enhanced in addition to operation effectiveness being maintained.

The presented AI-based CI/CD model is a basis of many improvements. It can offer higher levels of efficiency, more security, self-optimization, and intelligent decision support as the future implementations will offer the capacity to forecast, support multi-cloud, automated security and the whole adaptive deliveries of programs on its own. The ability to do this makes AI pipelines the cornerstone to the future of the next-generation enterprise DevOps practice.



V. CONCLUSION

The model of AI-based CI / CD pipeline is used in the research article and is implemented to maximize use of the secure.azure Kubernetes Services- NET applications. The chain of stages of the software delivery, monitoring the code commit and monitoring the latter after the implementation, etc. the framework will cover the issues of the main problems in the modern software delivery (how effective it will be at the moment of deployment, how stable it is going to be at its operation), and the specifics of its security. AI-supported predictive analytics, abnormalities, and adaptive orchestration of the data, decreasing human intervention, maximizing resource usage, and proactive vulnerability response are elements that facilitate fast and safe deployments.

Coverage of the framework analysis happened to be of great enhancement compared to the conventional CI/CD pipelines. It led to the 35 per cent of efficiency of deployment and 40 per cent decrease of the post deployment security incidents. The observation of mistakes performed automatically, intelligent rollback systems and utilization of the resources in the accompaniment of AI made the work true. It has already demonstrated to be scalable and agile to whatever workloads it is capable of supporting in reality and there is an indicator of flexibilities of the framework to the dynamic cloud-native business enterprise. The latter findings substantiate the assumption that not just the procedure of software simplification delivered is connected to the use of AI, but gives the scattered applications a more favorable security risk, overall.

The framework together with the already realized can be applied to improve in the future. They create additional opportunities to scale predictive maintenance, and adopt hybrid and multi-cloud applications, and do more to automatize security and constantly learn with self-optimization and AI-aided decision support to the DevOps teams. The improvements could also improve the deployment reliability, security and operational intelligence as well as cut on the expenditure and man-hours.

In conclusion, the proposed AI-based CI/CD model is a feasible, efficient and a solid solution to companies that want to embrace safe, efficient and functional deployed.Azure Kubernetes services app. Its architectural design, predictability, and in-built security setups, demonstrate how AI-influenced DevOps can transform cloud-native software delivery to provide a blueprint of smart, automorphic and secure deployments of programs in the modern enterprise.

REFERENCES

1. Octopus Deploy, "CI/CD Overview," Octopus.com. [Online]. Available: <https://octopus.com/devops/ci-cd>.
2. Microsoft Azure Blog, "Azure Pipelines is the CI/CD solution for any language, any platform, any cloud," Azure.Microsoft.com. [Online]. Available: <https://azure.microsoft.com/en-us/blog/azure-pipelines-is-the-ci-cd-solution-for-any-language-any-platform-any-cloud/>.
3. IBM, "CI/CD Pipeline," IBM.com. [Online]. Available: <https://www.ibm.com/think/topics/ci-cd-pipeline>.
4. Microsoft Learn, "Azure Kubernetes Service (AKS)," Learn.Microsoft.com. [Online]. Available: <https://learn.microsoft.com/en-us/azure/aks/>.
5. Microsoft Learn, "AKS Security Concepts," Learn.Microsoft.com. [Online]. Available: <https://learn.microsoft.com/en-us/azure/aks/concepts-security>.
6. Cloud Security Alliance, "The Evolution of DevSecOps with AI," CloudSecurityAlliance.org. [Online]. Available: <https://cloudsecurityalliance.org/blog/2024/11/22/the-evolution-of-devsecops-with-ai>.
7. OpenSSF, "ML Security in DevSecOps Whitepaper," OpenSSF.org. [Online]. Available: https://openssf.org/wp-content/uploads/2025/08/OpenSSF_MLSecOps_Whitepaper.pdf.
8. Checkmarx, "DevSecOps Best Practices in the Age of AI," Checkmarx.com. [Online]. Available: <https://checkmarx.com/learn/ai-security/devsecops-best-practices-in-the-age-of-ai/>.
9. Wikipedia, "DevOps Research and Assessment," Wikipedia.org. [Online]. Available: https://en.wikipedia.org/wiki/DevOps_Research_and_Assessment.
10. Wiz.io, "AKS Security Best Practices," Wiz.io. [Online]. Available: <https://www.wiz.io/academy/container-security/aks-security-best-practices>.