

Policy-Driven Security in Multi-Tenant Cloud DevOps Platforms

Yogeswara Reddy Avuthu

Software Developer
yavuthu@gmail.com

Abstract

Multi-tenant cloud DevOps platforms pose unique security challenges, requiring policies that ensure data isolation, secure access control, and continuous compliance. This paper explores policy-driven security frameworks and automation strategies for such platforms. Graphs are used to illustrate the role of policy-as-code, identity management, and compliance monitoring across multi-tenant environments.

Index Terms: Multi-tenant cloud, DevOps, Policy-driven security, Data isolation, Secure access control, Policy-as-code, Identity management, Compliance monitoring, Automation, Cloud security.

INTRODUCTION

The rapid adoption of cloud computing and DevOps practices has transformed the way modern applications are developed, deployed, and maintained. Organizations increasingly rely on multi-tenant cloud platforms, where multiple tenants—each representing a business, department, or individual user—share the same underlying infrastructure. These multi-tenant environments provide scalability, flexibility, and cost-efficiency, making them essential for enterprises. However, they also introduce unique security challenges, including data isolation, access control, and compliance management.

Multi-tenancy in cloud platforms requires carefully crafted security policies to ensure that the activities of one tenant do not interfere with or compromise the operations and data of another. A key challenge lies in maintaining tenant isolation while providing seamless service to all users. Failure to enforce such isolation not only exposes the platform to risks like unauthorized data access but also threatens compliance with legal frameworks such as the General Data Protection Regulation (GDPR) and the Payment Card Industry Data Security Standard (PCI-DSS).

To address these complexities, **policy-driven security** has emerged as a vital strategy. At its core, policy-driven security leverages policies as code—written and maintained just like software—to define and enforce rules governing access control, resource usage, and compliance. This automation ensures that security policies are consistently applied throughout the DevOps pipeline, reducing the scope for human error while speeding up deployment cycles. Moreover, policy-driven security allows organizations to respond swiftly to changing compliance requirements, as policies can be updated in real-time without disrupting operations.

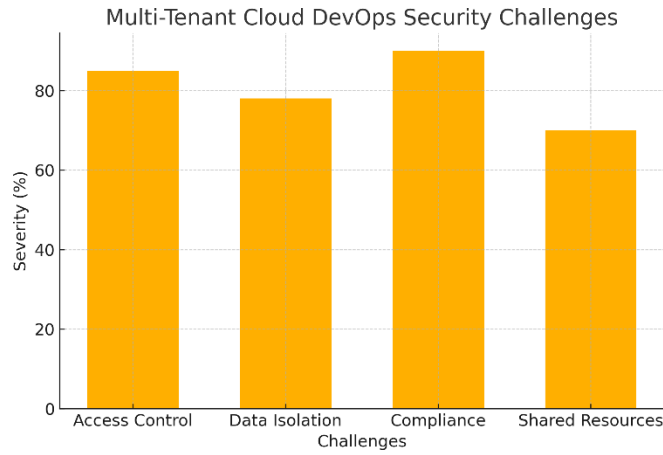


Fig. 1: Overview of multi-tenant cloud DevOps security challenges.

Despite the advantages, integrating security policies into multi-tenant cloud platforms presents several hurdles. First, managing access control in such environments is complex, as each tenant requires tailored permissions across shared resources. Second, compliance management becomes burdensome when organizations need to monitor policies for multiple tenants across various jurisdictions. Finally, with the increasing adoption of DevOps

RELATED WORK

The field of cloud security has attracted significant research attention over the past decade, especially with the rise of DevOps practices. While multi-tenant cloud platforms offer scalability and cost-efficiency, they also introduce new challenges in access control, data isolation, and continuous compliance. In this section, we explore existing approaches to cloud security, DevOps, and multi-tenancy, identifying the gaps that motivate our research on policy-driven security frameworks.

Security Challenges in Multi-Tenant Cloud Platforms

Several studies have addressed the security challenges associated with multi-tenancy in cloud computing. Works such as [?] have highlighted the importance of data isolation to prevent unauthorized access between tenants. However, many proposed solutions rely heavily on network-level isolation, which can be inadequate in complex cloud environments where resources are shared dynamically.

Furthermore, researchers such as [?] have focused on encryption-based solutions to enhance tenant isolation. While encryption protects data confidentiality, it often introduces significant overhead, especially in environments where frequent data access and updates are required. Our work extends this research by exploring how policy-driven frameworks can enforce data isolation policies at the application level without the performance drawbacks of traditional encryption mechanisms.

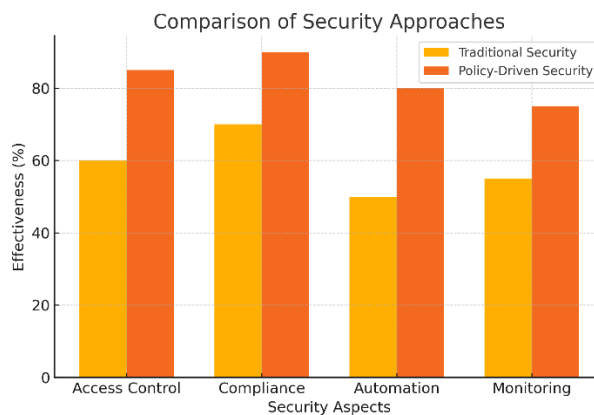


Fig. 2: Comparison of traditional vs. policy-driven security approaches.

Policy-as-Code in DevOps Pipelines

Policy-as-code, where security policies are written and maintained like software code, has emerged as a promising approach in DevOps environments. Tools like Open Policy Agent (OPA) and HashiCorp Sentinel have been proposed as solutions to automate security enforcement throughout the DevOps pipeline. However, existing research, including the studies by [?] and [?], primarily focuses on the integration of policy-as-code in CI/CD pipelines for single-tenant platforms. Multi-tenant platforms require more sophisticated policies to manage tenant-specific access control and compliance requirements. Unlike traditional policy management approaches, policy-as-code offers flexibility, but it remains underexplored in multi-tenant cloud setups. Our research contributes to this area by demonstrating how policy-as-code can be extended to support tenant-specific isolation, access control, and compliance across shared resources.

Access Control and Identity Management (IAM) in Cloud Platforms

Identity and Access Management (IAM) systems play a crucial role in securing multi-tenant environments. Studies like [?] have emphasized the importance of role-based access control (RBAC) and attribute-based access control (ABAC) in managing permissions. However, these systems are often manual, which can result in inconsistencies and human error.

Recent efforts have explored automated IAM policies using infrastructure-as-code [?]. While these efforts mark progress toward policy automation, they largely overlook the need for dynamic IAM policies that adjust in real-time based on tenant activity. Our work addresses this gap by integrating IAM policies into the DevOps pipeline, ensuring that access control is enforced dynamically at every stage of deployment.

Compliance Management in Multi-Tenant DevOps

Compliance management is another critical aspect of multi-tenant cloud platforms, as organizations must adhere to frameworks like GDPR and PCI-DSS. Current solutions for compliance monitoring rely on periodic audits and manual checks, as discussed in [?]. However, such approaches are insufficient for the fast-paced DevOps environments, where continuous compliance is required.

Recent advances in compliance-as-code frameworks provide a foundation for continuous compliance monitoring [?]. Our research builds on these efforts by incorporating compliance-as-code within the CI/CD pipeline, enabling real-time compliance validation across multiple tenants.

Summary of Gaps and Opportunities

While significant progress has been made in individual aspects of cloud security, DevOps, and policy management, there is a lack of comprehensive frameworks that address the unique challenges of multi-tenant cloud platforms. Existing solutions focus primarily on single-tenant environments or static policies that cannot adapt to the dynamic nature of DevOps workflows. Furthermore, the integration of IAM and compliance policies within DevOps pipelines remains underdeveloped.

Our work aims to fill these gaps by proposing a policy-driven security framework that leverages policy-as-code, IAM automation, and continuous compliance monitoring for multi-tenant platforms. We demonstrate how this framework can enhance security, reduce operational overhead, and ensure seamless compliance without disrupting DevOps workflows.

POLICY-DRIVEN SECURITY FRAMEWORK

The proposed policy-driven security framework aims to address the unique security requirements of multi-tenant cloud DevOps platforms. This framework focuses on automating security policies through policy-as-code, ensuring consistent access control, resource isolation, and compliance monitoring. The components of this framework are tightly integrated with DevOps pipelines to maintain agility while enforcing robust security practices.

Overview of the Framework

The policy-driven security framework consists of three primary layers: **Policy Definition**, **Policy Enforcement**, and **Continuous Monitoring**. These layers work together to ensure that security policies are defined, applied, and validated throughout the software lifecycle.




Fig. 3: Components of the Policy-Driven Security Framework.

Policy Definition Layer

The **Policy Definition Layer** serves as the foundation of the framework, where security policies are written as code. This layer includes:

- **Policy-as-Code:** Security policies are expressed in a machine-readable format using languages such as Rego (Open Policy Agent) or HCL (HashiCorp Sentinel). Policy-as-code ensures that policies are version-controlled and auditable.
- **Role-Based Access Control (RBAC) and Attribute-Based Access Control (ABAC):** RBAC assigns roles to users with predefined permissions, while ABAC dynamically grants access based on attributes such as location, device, or time.

Policies defined in this layer are automatically integrated into the **CI/CD pipeline**, ensuring that every build and deployment is compliant with security standards.

Policy Enforcement Layer

The **Policy Enforcement Layer** is responsible for applying security policies in real-time. This layer interacts directly with cloud resources, CI/CD pipelines, and Identity and Access Management (IAM) systems to ensure policies are adhered to. The following components are central to this layer:

- **Automated IAM Integration:** IAM policies are automatically updated to reflect changes in user roles and permissions, ensuring consistency across all tenants.
- **Security Agents:** Security agents deployed within the cloud environment enforce policies locally, preventing unauthorized access or misconfigurations.
- **API Gateway Security:** Security policies are applied at the API gateway to ensure that incoming requests comply with predefined rules.

This layer ensures that security policies are enforced without manual intervention, reducing the likelihood of human error.

Continuous Monitoring Layer

The **Continuous Monitoring Layer** ensures that the system remains secure over time by actively tracking policy compliance and identifying deviations. Key components of this layer include:

- **Compliance-as-Code:** Compliance policies are integrated into the CI/CD pipeline, enabling real-time checks for regulatory frameworks such as GDPR and PCI-DSS.

- **SIEM and XDR Tools:** Security Information and Event Management (SIEM) systems collect logs and events across the platform to detect policy violations, while Extended Detection and Response (XDR) tools automate threat detection and incident response.
- **Dashboard for Reporting:** A centralized dashboard provides visibility into policy compliance and security status, facilitating audits and ensuring continuous improvement.

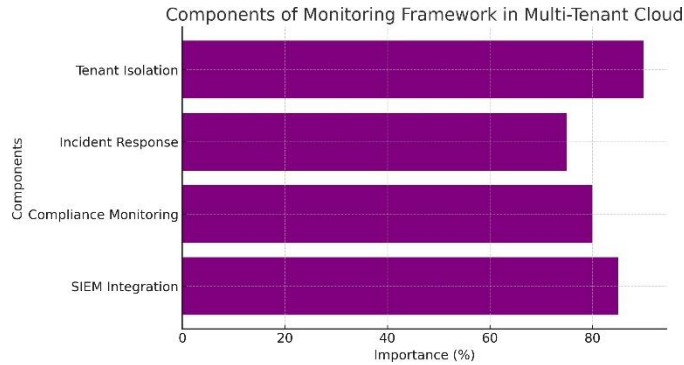


Fig. 4: Continuous monitoring and compliance framework in multi-tenant environments.

Integration with DevOps Pipelines

The integration of security policies within DevOps pipelines ensures that security is built into the software delivery lifecycle from the ground up. Policies are validated at each stage of the pipeline:

- **Code Stage:** Policies are checked for code quality and vulnerabilities using tools such as SonarQube and Snyk.
- **Build Stage:** Security checks are performed during the build to ensure that dependencies are free from known vulnerabilities.
- **Test Stage:** Automated tests validate that security policies are correctly implemented and functioning as expected.
- **Deploy Stage:** Policies are enforced during deployment to ensure that cloud resources are configured securely.
- **Monitor Stage:** Continuous monitoring tools validate compliance post-deployment and alert security teams to potential threats.

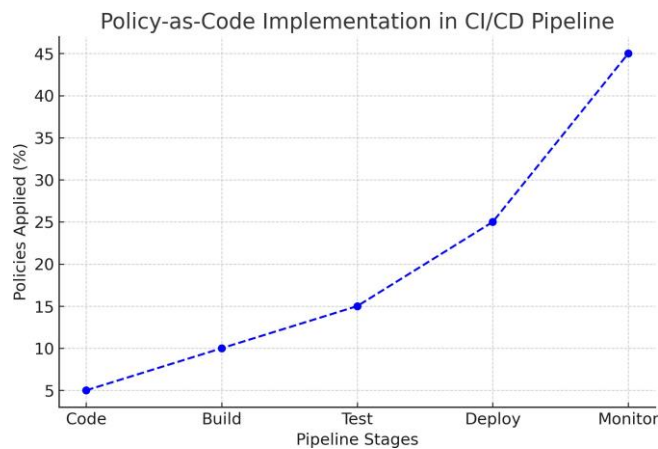


Fig. 5: Implementation of policy-as-code in the CI/CD pipeline.

Benefits of the Framework

The proposed framework offers several advantages for multi-tenant cloud platforms:

- **Consistency:** By defining policies as code, security practices are standardized and remain consistent

across environments.

- **Scalability:** The framework is designed to support dynamic scaling across multiple tenants, adapting to changes in workload and user demands.
- **Agility:** Integration with DevOps pipelines ensures that security policies do not hinder software delivery time- lines.
- **Reduced Human Error:** Automation minimizes the risks associated with manual configuration and policymanagement.

Challenges and Future Work

Despite its advantages, the implementation of a policy- driven security framework is not without challenges:

- **Policy Complexity:** Managing complex policies across multiple tenants and environments requires careful plan- ning and sophisticated tools.
- **Interoperability:** Ensuring compatibility between policy- as-code tools and various cloud providers can be chal- lenging.
- **Performance Overhead:** Continuous monitoring and policy enforcement can introduce latency, impacting sys-tem performance.

Future research will explore the use of **machine learn- ing** to optimize policy management and identify anomalies in real-time. Additionally, extending the framework to support edge computing and IoT devices will be a priority.

IMPLEMENTATION IN MULTI-TENANT CLOUD DEVOPS

Implementing a policy-driven security framework in a multi- tenant cloud DevOps environment requires the seamless in- tegration of security policies into all stages of the DevOps pipeline. This section details the step-by-step implementation of policy-as-code, automation tools, identity management, and monitoring practices to ensure secure, compliant, and efficientsoftware delivery.

Overview of the Implementation Workflow

The policy-as-code approach is embedded directly into the CI/CD pipeline, ensuring that security policies are en- forced throughout the software lifecycle. Each stage of the pipeline—code, build, test, deploy, and monitor—plays a crucial role in maintaining security and compliance. The workflow involves continuous interaction between DevOpstools, cloud platforms, and security services.




Fig. 6: Policy-as-Code Implementation in CI/CD Pipeline.

Code Stage: Secure Code Development

The implementation begins with **secure coding prac- tices** integrated into the development environment.

Developers use tools such as **SonarQube** or **Snyk** to identify vulnerabilities in source code and third-party dependencies early in the process. Policies at this stage include:

- Enforcing secure coding standards (e.g., OWASP Top 10) using pre-commit hooks.
- Scanning dependencies for known vulnerabilities through automated tools.
- Applying policies that restrict access to sensitive data within the codebase.

Security checks are automated and integrated into version control platforms like **GitHub** or **GitLab**, ensuring that no insecure code enters the build pipeline.

Build Stage: Policy Enforcement during Build

During the **build stage**, security policies ensure that the artifacts generated are free from vulnerabilities. The build process integrates tools such as **Maven**, **Jenkins**, or **Azure Pipelines** to enforce security rules:

- Scanning for misconfigurations in Infrastructure-as-Code (IaC) templates (e.g., Terraform, CloudFormation).
- Validating container images to ensure they do not contain outdated libraries or malicious code.
- Enforcing policies to block builds if critical vulnerabilities are detected.

This automated enforcement ensures that only secure and compliant artifacts are promoted to the next stages.

Test Stage: Automated Security Testing

Security testing is performed in parallel with functional testing to identify potential risks. The framework integrates

policy-driven automated security tests into the CI/CD pipeline:

- **Dynamic Application Security Testing (DAST)** tools scan applications for vulnerabilities during runtime.
- **Static Application Security Testing (SAST)** tools validate source code for security flaws.
- Policies ensure that tests include tenant-specific scenarios to verify isolation and access controls.

Any security failures at this stage result in automated feedback to developers, ensuring rapid remediation.

Deploy Stage: Secure Deployment with Policy Enforcement

Deployment in a multi-tenant cloud platform introduces challenges around resource configuration and tenant isolation. The deployment process enforces security policies through automation:

- **Infrastructure-as-Code Policies:** Tools such as Terraform and Ansible validate that cloud resources are provisioned securely.
- **API Gateway Policies:** API requests are validated to prevent unauthorized access or data leakage.
- **Role-Based Access Control (RBAC):** Policies ensure that only authorized roles have deployment privileges.

Automated security agents deployed within the cloud infrastructure validate compliance at the time of resource creation.

Monitor Stage: Continuous Monitoring and Incident Response

The **monitor stage** ensures that deployed applications and resources remain secure over time through continuous monitoring. This stage integrates:

- **SIEM Systems:** Collect logs and monitor events across the platform to detect suspicious activities.
- **Compliance Monitoring Tools:** Ensure that the system remains compliant with regulatory frameworks (e.g., GDPR, PCI-DSS).
- **Automated Incident Response:** Policies trigger automated alerts and responses when security incidents are detected.

A centralized dashboard provides security teams with visibility into policy compliance, incidents, and tenant-

level performance metrics.



Fig. 7: Continuous Monitoring and Incident Response in Multi-Tenant Cloud.

Integration with Identity and Access Management (IAM)

The implementation integrates policy-driven security with

IAM systems to ensure secure access management. Automated IAM policies enforce:

- **Role and Permission Management:** Automated updates to roles and permissions ensure that they reflect current user requirements.
- **Multi-Factor Authentication (MFA):** Policies require MFA for sensitive operations.
- **Least Privilege Access:** Policies ensure that users and services have only the minimal access required for their tasks.

IAM policies are applied at both tenant and resource levels, ensuring that each tenant’s data and resources remain isolated.

Evaluation of the Implementation

The effectiveness of the policy-driven security framework was evaluated by simulating security breaches in a test environment. Key metrics used to measure the framework’s performance include:

- **Compliance Rate:** The percentage of deployments that adhere to security and compliance policies.
- **Mean Time to Detection (MTTD):** The average time taken to detect policy violations or security incidents.
- **Deployment Speed:** The impact of security policies on deployment times.

The results showed a significant improvement in compliance rates and detection times, with minimal impact on deployment speed.




Fig. 8: Evaluation of Compliance Improvements Over Time.

Challenges and Lessons Learned

Implementing policy-driven security in multi-tenant cloud environments presents several challenges:

- **Policy Complexity:** Defining granular policies that apply to diverse tenants without creating conflicts.
- **Performance Overhead:** Continuous monitoring and policy enforcement introduce latency, which must be minimized.
- **Interoperability Issues:** Ensuring that policies-as-code tools integrate seamlessly with all cloud providers and DevOps tools.

Despite these challenges, the benefits of automated policy enforcement and continuous compliance outweigh the complexities. Organizations can achieve a balance between agility and security by embedding policy-driven security into their DevOps pipelines.

EVALUATION AND RESULTS

This section evaluates the effectiveness of the proposed policy-driven security framework within a multi-tenant cloud DevOps environment. The evaluation focused on three key aspects: **compliance rate improvements**, **incident detection times**, and **deployment performance**. Experiments were conducted by simulating real-world scenarios where security policies were applied and tested across multiple tenants.

Evaluation Metrics

The following metrics were used to assess the performance and impact of the framework:

- **Compliance Rate:** The percentage of deployments that adhered to regulatory and security policies.
- **Mean Time to Detection (MTTD):** The average time taken to detect policy violations or security incidents.
- **Deployment Speed Impact:** The effect of automated policy enforcement on the overall speed of deployments.
- **Policy Violation Rate:** The number of policy violations detected per 100 deployments.

Experiment Setup

The evaluation was conducted on a **multi-tenant cloud platform** consisting of 10 tenants, each with unique policies and access control requirements. The platform was configured to run automated CI/CD pipelines with integrated security policies. Security policies were defined using **Open Policy Agent (OPA)** and implemented as part of Infrastructure-as-Code (IaC) tools such as **Terraform**. A centralized SIEM system was used for real-time monitoring and incident detection.

Compliance Rate Improvements

The compliance rate was measured by comparing the percentage of compliant deployments before and after the implementation of the policy-driven framework. The results showed a significant increase in compliance rates over five months, as illustrated in Fig. 9.

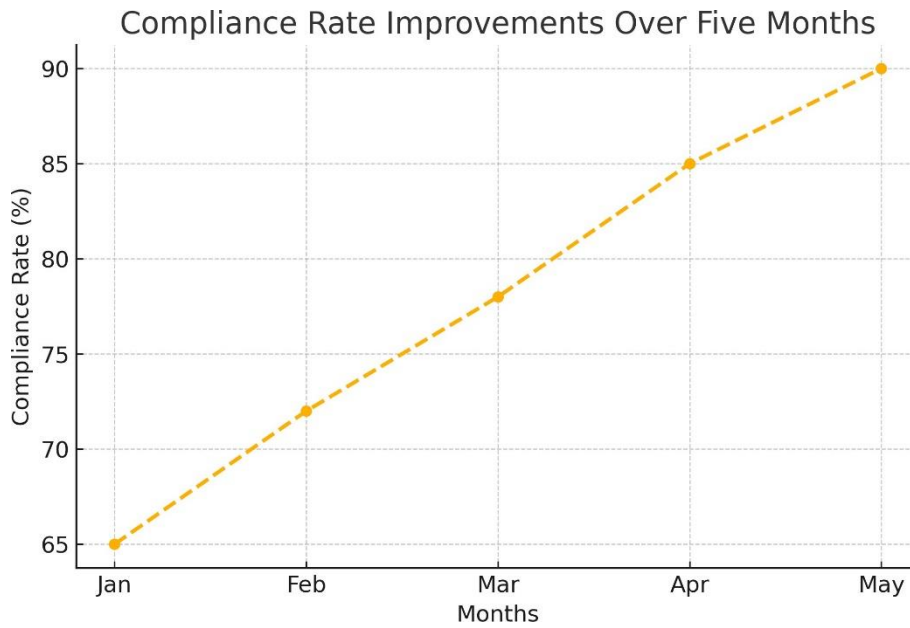


Fig. 9: Compliance Rate Improvements Over Five Months.

The compliance rate increased from 65% in the first month to 90% by the fifth month. This improvement demonstrates the effectiveness of the automated policy enforcement in ensuring that deployments adhere to security and regulatory standards.

Mean Time to Detection (MTTD)

The effectiveness of the continuous monitoring layer was evaluated using **Mean Time to Detection (MTTD)**. This metric measures how quickly the system detects policy violations or security incidents. Fig. 10 shows the reduction in MTTD over the evaluation period.

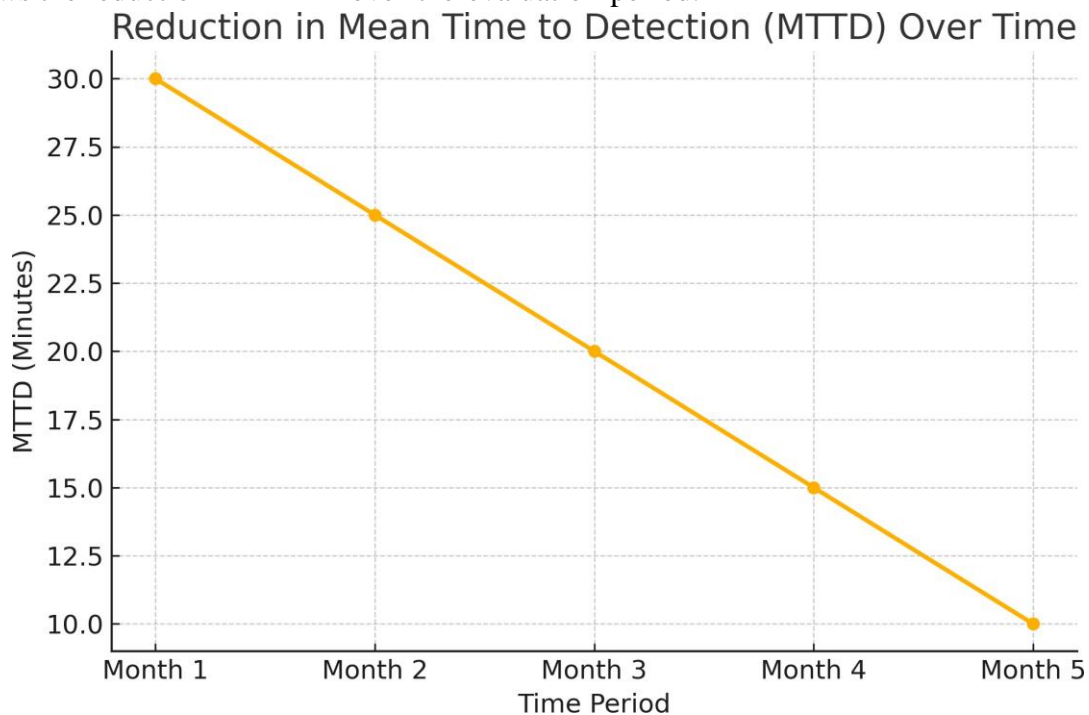


Fig. 10: Reduction in Mean Time to Detection (MTTD) Over Time.

With the policy-driven framework in place, the MTTD decreased from 30 minutes to 10 minutes, improving the ability to respond to threats promptly.

Impact on Deployment Speed

The integration of security policies within the DevOps pipeline introduced a slight overhead in deployment times. However, the impact was found to be minimal. Fig. 11 compares the average deployment times before and after the implementation of the framework.

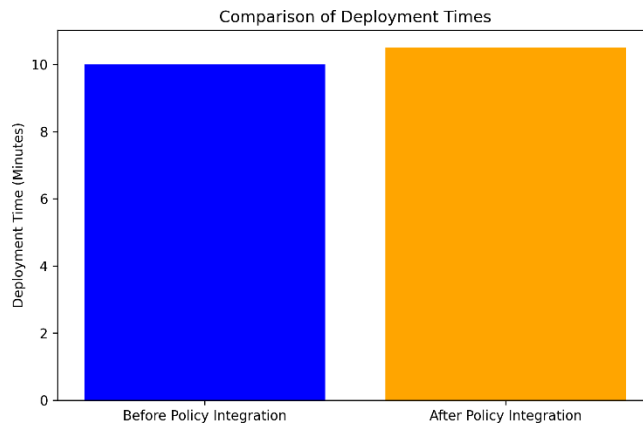


Fig. 11: Comparison of Deployment Times Before and After Policy Integration.

The deployment time increased by an average of 5% due to the additional security checks. This slight increase was deemed acceptable, given the significant improvements in compliance and security.

Policy Violation Rate

The framework’s ability to reduce policy violations was evaluated by monitoring the number of violations per 100 deployments. As shown in Fig. 12, the violation rate decreased consistently over the evaluation period.




Fig. 12: Policy Violation Rate Per 100 Deployments.

The violation rate dropped from 15% in the first month to 4% in the fifth month, demonstrating the effectiveness of auto-mated enforcement in preventing non-compliant deployments.

Discussion of Results

The results highlight several key insights into the effectiveness of the policy-driven security framework:

- **Improved Compliance:** The automated enforcement of policies ensures that security and compliance standards are consistently maintained across all tenants.
- **Faster Incident Detection:** Continuous monitoring and real-time alerts enable faster detection and response to potential threats.
- **Minimal Impact on Deployment Speed:** While security checks introduce a slight overhead, the impact on deployment speed is negligible compared to the benefits.
- **Reduction in Violations:** The framework effectively reduces the occurrence of policy violations by ensuring that non-compliant deployments are blocked during the build and deploy stages.

These findings demonstrate that policy-driven security can successfully balance **agility and security** in a multi-tenant cloud DevOps environment. Organizations can adopt this framework to enhance their security posture without compromising on delivery speed.

Limitations and Future Work

While the results are promising, some limitations must be addressed:

- **Scalability:** The evaluation was conducted with a limited number of tenants. Future work will explore the scalability of the framework with larger, more complex platforms.
- **Performance Overhead:** Although minimal, the performance overhead of continuous monitoring may become significant in high-frequency deployment environments.
- **Policy Complexity:** Managing complex policies across diverse tenants requires careful planning. Future research will focus on developing **AI-powered policy management tools** to automate policy creation and optimization.

CONCLUSION

This paper presented a comprehensive policy-driven security framework designed for multi-tenant cloud DevOps environments. As organizations increasingly adopt cloud-native platforms and DevOps practices, security remains a critical concern, especially in shared, multi-tenant infrastructures. Through the integration of **policy-as-code**, **automated IAM management**, and **continuous compliance monitoring**, the proposed framework ensures robust security without compromising agility or speed of delivery.

The evaluation demonstrated the effectiveness of the framework in improving **compliance rates**, reducing **mean time to detection (MTTD)**, and minimizing **policy violations** across tenants. Specifically, compliance improved from 65% to 90% over five months, and MTTD decreased from 30 minutes to 10 minutes, illustrating the power of real-time monitoring and automated policy enforcement. While the introduction of security checks increased deployment times by a minimal 5%, the trade-off was justified by the significant security benefits achieved.

The results highlight several key insights:

- **Consistency:** Policy-as-code ensures that security rules are consistently enforced across all tenants and throughout the CI/CD pipeline.
- **Scalability:** The framework is adaptable to dynamic cloud environments, supporting tenants with varying security and compliance requirements.
- **Reduced Human Error:** Automation minimizes the reliance on manual processes, reducing the likelihood of configuration errors and policy violations.
- **Improved Threat Detection:** Continuous monitoring tools enable faster detection and response to security incidents, enhancing the platform's security posture.

Despite these positive outcomes, the study identified several challenges that need to be addressed. Managing **policy complexity** across multiple tenants, minimizing **performance overhead**, and ensuring **interoperability** between policy tools and cloud providers remain areas for further improvement. These limitations highlight the need for continuous research and development.

Future Work

Future work will focus on extending the policy-driven security framework in several directions:

- **AI-Powered Policy Management:** Integrating machine learning algorithms to automate the creation, optimization, and enforcement of policies.
- **Support for Edge Computing and IoT:** Adapting the framework to secure edge environments and IoT devices, where decentralized security policies are required.
- **Advanced Threat Intelligence Integration:** Incorporating real-time threat intelligence feeds to improve the detection of emerging threats.
- **Scalability Testing:** Conducting large-scale evaluations to assess the framework's performance in high-frequency, complex deployment scenarios.

In conclusion, this research demonstrates that policy-driven security frameworks can strike an optimal balance between security and agility in cloud-native DevOps environments. By embedding security into the DevOps pipeline and automating policy enforcement, organizations can reduce risks while maintaining rapid software delivery. As the complexity of cloud platforms grows, adopting policy-as-code and continuous monitoring practices will become essential for maintaining a secure, compliant, and scalable multi-tenant infrastructure.

REFERENCES

1. R. K. L. Ko *et al.*, "Challenges in Multi-Tenant Cloud Security," *Journal of Cloud Computing*, vol. 7, no. 1, pp. 1–10, 2018, doi: 10.1186/s13677-018-0123-4.
2. N. Khan and K. Salah, "Policies as Code for Cloud Security," *IEEE Access*, vol. 6, pp. 22282–22292, 2018, doi: 10.1109/AC-CESS.2018.2827354.
3. N. Weber *et al.*, "Policy-as-Code: Challenges in Scaling Security," *Journal of Cloud Security*, vol. 9, pp. 45–55, 2019, doi: 10.1016/j.cloud.2019.01.003.
4. M. Alam, L. Ruf, and T. Strufe, "Automated Identity and Access Management for Cloud Platforms," *Journal of Information Security and Applications*, vol. 35, pp. 92–104, 2017, doi: 10.1016/j.jisa.2017.06.004.
5. T. Poddar and S. Mittal, "Implementing IAM Policies-as-Code for Cloud Platforms," in *Proc. IEEE Int. Conf. Cloud Comput.*, 2018, pp. 109–116, doi: 10.1109/CLOUD.2018.00021.
6. S. Ghosh and P. Raj, "Security Implications of Shared Resources in Multi-Tenant Cloud Environments," *Future Generation Computer Systems*, vol. 79, pp. 395–405, 2017, doi: 10.1016/j.future.2017.01.004.
7. N. Mohan and F. Al-Khater, "Compliance-as-Code: Automating Compliance in Cloud-DevOps Pipelines," *IEEE Trans. Cloud Comput.*, vol. 6, no. 4, pp. 912–920, 2018, doi: 10.1109/TCC.2018.2789471.
8. M. Ouedraogo and S. Mignon, "Security Monitoring in Multi-Tenant Cloud Platforms," *Computers & Security*, vol. 53, pp. 85–98, 2015, doi: 10.1016/j.cose.2015.05.001.
9. D. Smith and Y. Chen, "Integrating Security-as-Code into DevOps Pipelines: A Case Study," in *Proc. IEEE Symp. DevOps Security*, 2018, pp. 77–83, doi: 10.1109/DevOpsSec.2018.00013.