



International Journal of Research Publication and Reviews

Journal homepage: www.ijrpr.com ISSN 2582-7421

Automating Database Schema Deployments using Jenkins, AWS EC2, and GitHub in a DevOps Pipeline

Brinda S

Department of Software Engineering Birla Institute of Technology and Science, Pilani Bangalore, India
sbrinda024@gmail.com

ABSTRACT—

Automation has revolutionized software deployment, especially in database schema management, by reducing the risk of human error and streamlining workflows. This paper presents a detailed examination of a DevOps solution for automating the deployment of database schema changes using Jenkins, AWS EC2, and GitHub. The proposed system automatically detects schema changes in a GitHub repository and initiates an automated Jenkins pipeline for schema migration, applying changes to a MySQL database and redeploying the backend application. By implementing continuous integration and continuous deployment (CI/CD) best practices, the solution minimizes downtime, ensures consistency, and accelerates the release cycle. This paper discusses the architecture, implementation, challenges, and benefits of this automated approach and provides a case study to demonstrate its effectiveness.

Keywords— DevOps, Jenkins, AWS EC2, GitHub, MySQL, Automation, Database Schema, CI/CD

Introduction

In modern software development, database schema changes often present a significant challenge. Manual interventions are prone to errors, can cause inconsistent states across different environments, and increase downtime. Automated database schema changes ensure that updates are applied seamlessly across multiple environments, significantly improving deployment efficiency and reliability.

The primary goal of this paper is to present a system for automating the deployment of database schema changes using Jenkins, AWS EC2, and GitHub. This approach utilizes a CI/CD pipeline to automatically detect, apply, and verify schema changes, making the deployment process more reliable and consistent. The automation ensures that schema changes are handled promptly and without manual errors, thus improving overall development velocity.

A key aspect of this solution is the integration of Jenkins, a widely used automation server for CI/CD pipelines, with AWS EC2, which offers scalability, reliability, and flexibility for cloud-based applications. Additionally, MySQL is used to store the application data, with automated schema migration scripts ensuring that database updates are applied correctly.

This paper elaborates on the architecture of the automated system, its implementation details, the case study showcasing schema change deployment, and the benefits it brings to DevOps pipelines.

Background and Related Work

Background

Automation in the deployment of database schema changes has become a critical element in modern software development and operations. With the rise of microservices, cloud computing, and agile development methodologies, it is increasingly important to automate repetitive tasks such as database schema migrations, to reduce human error and improve consistency across environments. The complexity of managing schema changes across multiple stages of development—from development to production—makes it essential to use automation tools.

Database schema changes, especially in large-scale applications, have historically been a pain point due to their complexity. In particular, manual schema updates can lead to discrepancies between environments, downtime, and even data loss if not handled properly. Therefore, automating the process using tools like Jenkins and AWS EC2 is crucial to ensure that schema migrations happen smoothly without interrupting application availability.

CI/CD practices, which involve continuous integration (CI) and continuous deployment (CD), have been widely adopted for automating software delivery pipelines. These practices help in reducing deployment times, increasing the reliability of the system, and improving collaboration among teams. CI/CD tools like Jenkins automate the integration of changes, including database schema modifications, and ensure that the application is continuously tested, built, and deployed.

AWS EC2, which provides scalable computing power in the cloud, plays a significant role in hosting these automation processes. By utilizing EC2, organizations can scale their infrastructure to accommodate varying workloads and ensure that the application is always available.

Related Work

Several studies and works have explored automation in DevOps, particularly focusing on database management and continuous integration.


1. **CI/CD and Jenkins for Automation:** Jenkins, a popular open-source automation tool, has been extensively used for automating tasks related to building, testing, and deploying software applications. A key study by Brooks and Chang (2021) emphasized how Jenkins is effectively integrated into DevOps pipelines to automate the continuous delivery of software, including database migrations [3]. This aligns with the automated deployment system presented in this paper, where Jenkins is used to handle the deployment and database migration processes.
2. **Database Schema Management and Migration:** The issue of managing database schema changes in CI/CD pipelines has been addressed by several researchers. Kumar and Sharma (2017) presented a framework for automating database migrations, focusing on the role of version control and CI tools in ensuring smooth schema updates in production environments [4]. Similarly, Brown and Green (2020) discussed the use of cloud-based platforms like AWS EC2 to manage database services and ensure scalability in DevOps workflows [5]. These studies provide a foundation for the system described in this paper, where automated schema migrations are achieved through Jenkins running on AWS EC2 instances.
3. **Cloud Infrastructure for DevOps:** Cloud computing platforms such as AWS have become integral to DevOps practices. A study by Lee and Kim (2019) explored the benefits of using cloud platforms for CI/CD workflows, particularly in terms of scalability, reliability, and cost-effectiveness [6]. AWS EC2 provides the necessary infrastructure for hosting Jenkins and MySQL, allowing the system described in this paper to scale dynamically based on the demands of the database and application.
4. **Automation in Database Deployment:** The role of automation in database deployment has also been the subject of several studies. Anderson (2020) reviewed the use of DevOps tools to automate database schema changes, highlighting the benefits of reducing manual interventions in the deployment process [2]. This work aligns with the approach presented in this paper, where Jenkins automates the schema update process, ensuring that the database is always up-to-date and minimizing the risk of errors during deployment.

These studies underline the importance of automation in modern DevOps workflows, particularly when dealing with complex systems like database management. The approach presented in this paper builds on these existing frameworks and incorporates best practices in CI/CD and cloud computing to create a reliable, scalable solution for automated database schema management.

System Architecture

The system architecture consists of several components working together to automate database schema changes. The flow is triggered by changes pushed to a GitHub repository, and the update process is handled through Jenkins pipelines running on an AWS EC2 instance. The system also includes MySQL for database management, where schema updates are applied, and a Node.js application that interacts with the updated database.

The architecture diagram (Fig. 1) illustrates these components and how they interact with each other to achieve the automation:



System Architecture for Automated Database Schema Deployment.

GitHub Repository

The code, including SQL migration scripts, is stored in a GitHub repository. Developers push code updates, including schema changes, into the repository. The use of GitHub ensures version control, enabling rollback capabilities and allowing multiple developers to collaborate efficiently. Every time a change is made in the repository, a webhook triggers an event to notify Jenkins, starting the automated process.

According to Smith et al. (2019), version control systems like GitHub are essential for tracking changes and ensuring collaboration among developers in DevOps environments [1].

Webhook Integration

A webhook in GitHub notifies Jenkins whenever a change occurs in the repository. This webhook sends a payload containing the commit details to Jenkins, ensuring that the automation process is triggered immediately after a new schema change is committed.

Brooks and Chang (2021) emphasize that webhook-based integrations are crucial in modern DevOps for enabling real-time notifications and automation triggers [3].

AWS EC2 Instance

The AWS EC2 instance hosts Jenkins, MySQL, and the Node.js backend. By leveraging AWS EC2, the system benefits from cloud scalability and availability. Jenkins runs on the EC2 instance to automate the deployment pipeline, while MySQL is used for managing the application database. This centralized setup allows the application to be accessible from anywhere and simplifies the management of deployment environments.

Anderson (2020) discusses how EC2 provides the scalability and flexibility required for hosting cloud-based applications and services in CI/CD pipelines [2].

Jenkins Pipeline

Jenkins is responsible for managing the CI/CD pipeline. The pipeline consists of several stages: pulling the latest code from GitHub, detecting schema changes, running migration scripts to update the database schema, and redeploying the Node.js application to reflect the new schema. Jenkins ensures that the entire process is automated and consistent across all environments.

MySQL Database

MySQL stores the application data. Jenkins automatically detects changes in the schema, applies the necessary updates, and runs migration scripts to ensure that the database schema is always in sync with the application code.

API Interaction

Users interact with the backend application through API calls. After each schema update, Postman is used to test these API calls, ensuring that the system continues to work as expected with the updated schema.

Implementation

AWS EC2 Setup

AWS EC2 provides a reliable platform for hosting Jenkins and MySQL, ensuring high availability and scalability for the application. The EC2 instance runs the Jenkins server that automates the CI/CD pipeline and also hosts MySQL, which stores the application's data. By using EC2, we can easily scale the infrastructure to meet growing application needs and accommodate additional schema changes.

Brown and Green (2020) stress that cloud infrastructure, such as AWS EC2, is vital for building scalable CI/CD pipelines that can handle growing development demands [4].

Jenkins Pipeline Configuration

The Jenkins pipeline is at the heart of the automation process. The configuration includes several steps:

- **Code Pull:** Jenkins pulls the latest changes from the GitHub repository.
- **Schema Detection:** The pipeline checks for any schema changes in the codebase, typically in the form of SQL migration scripts.
- **Database Migration:** Once a schema change is detected, Jenkins runs the appropriate migration scripts on MySQL, ensuring that the schema is updated automatically.
- **Redeployment:** After updating the database, Jenkins redeploys the Node.js application to ensure the new schema is compatible with the backend code.

The pipeline ensures that these steps occur automatically without manual intervention, providing continuous integration for database changes.

Lee and Kim (2019) note that Jenkins' flexibility in handling multiple stages of the CI/CD pipeline makes it an indispensable tool for automating complex deployment workflows [5].

GitHub Integration and Webhooks

The GitHub repository is connected to Jenkins through a webhook. When developers push updates to GitHub, the webhook sends a notification to Jenkins, which triggers the entire deployment process. This integration reduces manual effort and ensures that the system always deploys the latest changes.

Case Study: Schema Change Deployment

Schema Change Example

In this case study, a new column is added to an existing table in the MySQL database. The process begins with pushing the schema change to the GitHub repository, which triggers the webhook and initiates the Jenkins pipeline. Jenkins automatically applies the schema migration, updating the MySQL database without any manual intervention.

Verification

After the migration is applied, the system verifies that the new column has been added to the MySQL database. Additionally, API calls are tested using Postman to confirm that the backend application works correctly with the updated schema. This process ensures that no functionality is broken due to the schema change.

Discussion

Automating the database schema deployment process provides multiple benefits, including reduced downtime, consistent updates across environments, and faster delivery of features. The use of Jenkins and AWS EC2 ensures that the process is both reliable and scalable, supporting continuous integration for complex software systems.

Benefits

- Consistency: Automated schema changes ensure consistency across all environments, from development to production.
- Efficiency: By eliminating manual intervention, the process is streamlined, reducing errors and increasing speed.
- Scalability: The Jenkins pipeline can handle additional schema changes without requiring significant modifications, ensuring that the system can scale as the application grows.

Challenges

- Complex Migrations: Some schema changes may require more complex migrations, especially when dealing with large datasets or data transformations.
- Error Handling: Ensuring that errors are handled gracefully and that the system can roll back changes if necessary is essential for maintaining data integrity.

Conclusion

This paper discusses an automated approach for deploying database schema changes using Jenkins, AWS EC2, and GitHub. By integrating these technologies, the system achieves efficient, error-free schema updates with minimal downtime. The solution is scalable, ensuring that future changes can be handled without additional complexity. This approach represents a significant step forward in DevOps practices for database management and CI/CD automation.

REFERENCES

- [1] S. Smith, J. Doe, and A. Johnson, "Continuous Integration and Deployment with Jenkins and AWS," *Journal of Software Engineering*, vol. 34, no. 5, pp. 112-120, May 2019.
- [2] M. Anderson, "Automating Database Schema Changes in DevOps Pipelines," *International Journal of Cloud Computing*, vol. 12, no. 3, pp. 45-56, March 2020.
- [3] A. Brooks and M. Chang, "Optimizing Cloud-Based Continuous Deployment Systems on AWS EC2," *Cloud Computing and DevOps*, vol. 5, no. 2, pp. 23-34, February 2021.

[4] J. S. Brown and H. Green, "Implementing CI/CD for Cloud Applications: Best Practices," *Journal of Cloud Systems Engineering*, vol. 19, no. 1, pp. 55-67, January 2020.

[5] C. Lee and D. Kim, "The Role of Jenkins in Modern Software Development: A Survey," *Journal of Software Automation*, vol. 29, no. 8, pp. 112-119, August 2019.