RelengDesk: An Enterprise Grade Release Engineering Monitoring and Analytics System

Saumya Shrivastava*, Satyendra Tiwari**

Adobe Systems, Noida

India

ABSTRACT

Modern software development environments face significant challenges in release engineering, particularly in managing complex build processes, ensuring real-time status tracking, and generating actionable analytics. This paper presents RelengDesk, a Spring Boot-based release engineering system that combines a microservices architecture with event-driven processing using Apache Kafka and flexible data persistence with MongoDB. RelengDesk enables real-time build status updates, comprehensive analytics, and dynamic channel selection for flexible event handling. The system's platform-agnostic architecture allows integration with any CI/CD pipeline, providing centralized access control and customizable views per product, while eliminating the limitations of traditional plugin-based solutions. By centralizing build tracking and analytics, RelengDesk reduces the time required for teams to monitor and verify build statuses, accelerating feedback cycles and enhancing productivity. The system's extensible design and robust error handling make it a strong candidate for enterprise deployment in continuous integration and delivery environments.

Keywords :- Engineering, Build Tracking, Microservices Architecture, Event Processing, Real-Time Analytics

I. INTRODUCTION

Modern software development practices have introduced significant challenges for release engineering [1]. Organizations must efficiently manage multiple concurrent builds, provide real-time status updates, and generate meaningful analytics [2]. The need for rapid deployment and continuous integration has created a demand for sophisticated build management solutions [3]. Additionally, maintaining data consistency across distributed systems and ensuring reliable error handling are now critical concerns in enterprise environments. The scale of operations in large organizations, where hundreds or thousands of builds may be processed simultaneously, further exacerbates these challenges [4].

Traditional monitoring solutions, particularly those tied to specific CI/CD platforms, present several limitations. These include scattered job information across multiple plugins, limited visibility into parallel builds within the same job, and performance degradation due to API-based data fetching. Furthermore, these solutions often struggle with multi-master node architectures and lack the flexibility to adapt to different CI/CD platforms. The reliance on platform-specific plugins and APIs creates vendor lock-in, making it difficult for organizations to switch between different CI/CD tools or maintain multiple platforms simultaneously.

To address these issues, this research presents RelengDesk, a comprehensive release engineering system designed to efficiently handle multiple concurrent builds while providing real-time status updates and analytics. The system emphasizes robust error handling and logging to ensure reliable operation under heavy load [5]. Furthermore, RelengDesk is built to maintain data consistency across distributed systems and scale effectively for enterprise-level build volumes [6]. Its platformagnostic architecture allows integration with any CI/CD pipeline, not just Jenkins, providing a unified interface for monitoring builds across multiple platforms.

The architecture of RelengDesk is designed to be flexible and extensible, allowing seamless integration with existing development tools and workflows. Its analytics capabilities provide insights into build processes, helping organizations identify and resolve bottlenecks in their development pipelines [7]. By focusing on scalability, reliability, and performance, RelengDesk aims to set new standards for release engineering systems in enterprise environments. The system's centralized access control mechanism eliminates the need for individual platform permissions management, significantly simplifying the administration of build access across multiple instances.

In summary, this research contributes a robust, scalable, and efficient solution for release engineering, supporting organizations in achieving faster and more reliable software delivery. RelengDesk's comprehensive approach to build management, real-time analytics, and reliable error handling makes it a valuable tool for improving release engineering processes. The system's platform-agnostic design and centralized access control provide a significant advantage over traditional solutions, enabling organizations to maintain flexibility in their choice of CI/CD tools while ensuring efficient build monitoring and management.

II. LITERATURE REVIEW

A. Comparative Analysis of Existing Solutions

The evolution of release engineering solutions has been marked by the emergence of both traditional and modern tools,

International Journal of Information Technology (IJIT) - Volume 11 Issue 3, May - Jun 2025

each with distinct strengths and limitations [8]. Traditional build systems such as Jenkins, Bamboo, and TeamCity have provided foundational support for build automation and continuous integration. However, these platforms often encounter scalability challenges and offer limited analytics, making them less suitable for the demands of contemporary enterprise environments [9]. Modern CI/CD platforms, including GitHub Actions, GitLab CI, and CircleCI, have introduced advanced features such as cloud-native integration and seamless version control connectivity [10]. Despite these enhancements, organizations frequently face issues related to vendor lock-in and restricted customization, which can hinder adaptation to specific infrastructure requirements.

General-purpose monitoring and analytics tools like Splunk, Grafana, and Kibana are widely used for build tracking and process monitoring. While these platforms offer robust data visualization and analysis capabilities, their application to release engineering typically requires significant customization and configuration. For instance, Splunk excels in log analysis but demands extensive dashboard and alert setup for build tracking. Grafana provides powerful timeseries visualization but lacks native support for build process monitoring. Kibana, though effective for log visualization, necessitates Elasticsearch integration and considerable configuration for build-specific use cases.

A comprehensive analysis of existing solutions reveals distinct trade-offs between functionality, setup effort, and cost. Table I presents a comparison of prominent monitoring solutions available in the market:

TABLE I
COMPARISON OF EXISTING MONITORING SOLUTIONS

Solution	Aggregates Multiple Masters?	Setup Effort	Advantages	Limitations
Prometheus + Grafana	Yes	Medium	 Best for metrics & visual dashboards Highly customizable Open source 	 Requires extensive setup for build monitoring No native build process tracking High learning curve Additional configuration needed for build-specific metrics
Jenkins Operations Center	Yes	Low (but Paid)	 Enterprise-level control Easy setup Built-in security 	 High licensing costs Difficult to switch to other solutions (vendor lock-in) Limited customization Additional costs for scaling
RelengDesk (Custom Dashboard)	Yes	Medium	 Platform-agnostic architecture Real-time build tracking Centralized access control Customizable views per product Parallel build monitoring Selective build display No impact on Jenkins master performance 	 Initial setup required Custom build service integration needed Requires maintenance
Build Monitor View Plugin	No	Low	 Visual job dashboard Easy to install Basic monitoring 	 Limited to single Jenkins instance No multi-master support Basic features only No advanced analytics
Jenkins API	Yes	Medium	 Direct access to Jenkins data Flexible integration Real-time data 	 Performance impact on Jenkins master API rate limiting Requires custom development No built-in visualization

Existing monitoring solutions, particularly Jenkins plugins, present several limitations that impact their effectiveness in enterprise environments. These limitations include scattered job information across multiple plugins, limited visibility into parallel builds within the same job, inability to selectively display build information, and performance degradation due to API-based data fetching. Furthermore, these solutions often struggle with multi-master node architectures and lack the flexibility to adapt to different CI/CD platforms.

B. Technology Stack Analysis

1) *Spring Boot Framework:* Spring Boot serves as the core framework for RelengDesk, offering production-ready features such as embedded servers, security, and monitoring capabilities [11]. Its robust dependency injection and strong

International Journal of Information Technology (IJIT) - Volume 11 Issue 3, May - Jun 2025

typing ensure maintainable and reliable code, while its performance in handling concurrent requests makes it ideal for high throughput build processing. The mature Spring ecosystem and comprehensive documentation further enhance its suitability for enterprise applications.

2) *MongoDB Database:* MongoDB is selected as the primary data store to accommodate the flexible and diverse nature of build artifacts and metadata [12]. Its schema-less design supports efficient storage and retrieval of varied data types, while high performance and horizontal scalability enable the management of large build volumes and concurrent operations. MongoDB's advanced query capabilities facilitate complex analytics and real-time status updates, and its replication and sharding features ensure data reliability in distributed environments.

3) *Apache Kafka:* Apache Kafka is integrated to provide robust event processing, essential for modern release engineering systems [13]. Kafka's high throughput and fault tolerance support efficient processing and reliable delivery of build events, even during system failures. Its scalability and message persistence features ensure that no build events are lost, maintaining system integrity during restarts or outages.

The implementation of RelengDesk is guided by key software engineering principles to ensure robustness, scalability, and maintainability [15]. Modularity is achieved by defining clear interfaces for each service domain and implementing them

C. Comparison with Alternative Technologies

The technical implementation of RelengDesk demonstrates significant advantages over alternative monitoring solutions in several key areas. Traditional solutions often rely on direct API calls to Jenkins master, causing performance degradation during high load scenarios [8]. In contrast, RelengDesk's event-driven architecture using Kafka ensures efficient event processing without impacting build performance. The system's microservices architecture enables independent scaling of components based on load, providing better resource utilization and system stability.

In terms of data management and analytics, RelengDesk's MongoDB-based data model is optimized for build tracking and analytics, providing real-time insights without additional data processing overhead. This optimization enables faster query response times and more efficient resource utilization, particularly in environments with large numbers of concurrent builds. The system's architecture eliminates the need for complex data transformations required by general-purpose tools, resulting in more efficient and reliable build monitoring.

The integration and extensibility of RelengDesk sets it apart from traditional solutions that are tightly coupled with specific CI/CD platforms [14]. The system's platform-agnostic design allows integration with any CI/CD pipeline, providing organizations with the flexibility to use their preferred tools while maintaining centralized monitoring and control. The modular architecture enables easy addition of new features and integrations, ensuring the system can evolve with changing requirements and technologies.

Performance and resource utilization are critical factors in enterprise environments, where traditional API-based solutions often impact Jenkins master performance during high load. RelengDesk's asynchronous event processing ensures minimal impact on build systems, while its efficient resource utilization enables handling of large-scale build operations [4]. This approach results in more reliable build monitoring and better overall system performance.

Security and access control represent another area where RelengDesk provides significant advantages. Traditional solutions require separate permission management for each platform, increasing administrative overhead and potential security risks [7]. RelengDesk's centralized access control provides unified authentication and authorization, implementing enterprise-grade security features while maintaining flexibility. This centralized approach simplifies security management and ensures consistent access control across all integrated systems.

These technical advantages position RelengDesk as a superior solution for enterprise release engineering environments, particularly in scenarios requiring high scalability, real-time monitoring, and platform flexibility [10]. The system's architecture and implementation choices address the limitations of existing solutions while providing enhanced functionality and performance. By focusing on these key technical aspects, RelengDesk delivers a more robust, efficient, and maintainable solution for modern release engineering challenges.

III. METHODOLOGY

A. System Architecture

The architecture of RelengDesk is designed to centralize and streamline the release engineering process by integrating build tracking, event processing, and analytics within a modular, microservices-based system [14]. As shown in Fig. 1, the platform is organized around several core services that interact through well-defined channels and persistent data stores. This modular separation ensures that each component can be developed, deployed, and scaled independently, supporting both reliability and extensibility [1].

A key architectural aspect of the system is its layered design and dynamic channel selection mechanism. The system is organized into distinct layers: External Layer, API Layer, Service Layer, Data Layer, and Database Layer. The Build Service, comprising multiple builder instances, determines the appropriate communication channel—either synchronous API or asynchronous Kafka—based on the current system configuration, which is stored and retrieved from a local cache. This design allows the system to switch to Kafka-based processing under high request volumes for scalability and stability, or to use the API channel for lower volumes and immediate processing. Build events are routed accordingly, ensuring efficient and flexible event handling. The RelengDesk Service acts as the central coordinator, processing incoming events, updating MongoDB for persistent storage, and providing real-time status updates to the dashboard. The service also exposes an endpoint for dynamically retrieving the current channel configuration, further enhancing system adaptability and operational flexibility.

B. Implementation Approach

The implementation of RelengDesk adheres to key software engineering principles to ensure robustness, scalability, and maintainability [15]. Modularity is achieved by defining clear interfaces for each service domain and implementing them in separate classes, enabling independent development, testing, and deployment. For example, build tracking and analytics services can be scaled or updated independently, minimizing system-wide risks.

Scalability is addressed through stateless service implementations, allowing horizontal scaling via multiple service instances. Load balancing is managed by the underlying infrastructure, and efficient resource utilization is achieved through careful management of dependencies and data access patterns. Kafka further enhances scalability by decoupling event producers and consumers.

Reliability is ensured through comprehensive error handling and data consistency mechanisms. Aspect-Oriented Programming (AOP) is used for cross-cutting concerns such as exception handling and logging, ensuring consistent error management. Data consistency is maintained through transactional repository operations and robust event processing. Fault tolerance is achieved via Kafka's message persistence and MongoDB's replication and sharding features.

The system is designed for extensibility, with interfaces and dependency injection facilitating the integration of new features or the replacement of components. Additional analytics modules or alternative data stores can be incorporated with minimal changes to the codebase.

C. Architectural and Workflow Diagrams

To provide a comprehensive understanding of the system's architecture and workflow, a few UML diagrams are presented in this section. Each diagram highlights a different aspect of the RelengDesk platform, including high-level design, core service interactions, data flow, and channel selection mechanisms. These diagrams illustrate the modular structure, component relationships, and key processes underpinning system functionality, facilitating a clearer explanation of design decisions and operational principles.







Fig. 2 Depiction of the flow of build data from external sources (API/Kafka) through the service layer (BuildTrackingService, AnalyticsService, KafkaConsumerService) to MongoDB storage via repository interfaces.

International Journal of Information Technology (IJIT) – Volume 11 Issue 3, May - Jun 2025



Fig. 3 Illustration of the modular architecture of RelengDesk, which monitors and analyzes build artifacts from Jenkins CI/CD pipelines across multiple platforms. It shows the integration of build status tracking and analytics modules, with service interfaces and implementations managing the build lifecycle, real-time monitoring, and analytics. The clear relationships between core entities highlight the system's extensibility and efficient coordination between tracking and analytics functionalities.

These diagrams illustrate the modular structure, component relationships, and key processes underpinning system functionality, facilitating a clearer explanation of design decisions and operational principles.

IV. IMPLEMENTATION DETAILS

A. System Architecture and Core Components

The implementation of RelengDesk is structured around a set of core components, each responsible for a distinct aspect

of the release engineering workflow [9]. The primary components include the Build Service, RelengDesk Service, Analytics Service, and Event Processing Service. The Build Service is designed to be platform-agnostic, supporting integration with various CI/CD pipelines such as Jenkins, CircleCI, and others. This is achieved through a cluster of builder instances and custom scripts that push build events to the backend based on stage and job-level notifications.

The RelengDesk Service acts as the central coordinator, receiving build events via either synchronous API calls or asynchronous Kafka messages [12], as determined by the dynamic channel selection mechanism. This service processes incoming events, updates build statuses and orchestrates the flow of information throughout the system. It also manages persistent storage in MongoDB and provides real-time status updates to the dashboard for user visibility. The Analytics Service operates independently from the core build tracking logic, aggregating build metrics, tracking failures, and analyzing performance characteristics. This separation allows for independent scaling and evolution of analytics capabilities, ensuring that performance analysis does not impact core build management. The Event Processing Service, implemented using Apache Kafka, enables asynchronous processing of build events, ensuring reliable event handling even under heavy load or transient failures.

B. Data Management and Service Layer

The data models in RelengDesk are designed to capture essential information for effective build tracking and analytics [6]. The primary data entities are the Build and BuildAnalytics classes, each annotated as a persistent entity for storage in MongoDB. The Build entity encapsulates all relevant information about a software build, including a unique identifier, current status, timestamps for the start and end of the build process, a list of generated artifacts, and a flexible metadata map for additional build-related information. This design accommodates a wide range of build scenarios, from simple single-step builds to complex, multi-stage pipelines, and allows the system to adapt to evolving requirements without changes to the underlying schema.

The service layer in RelengDesk is implemented using the Spring Framework, leveraging its dependency injection and component management features [11]. Each core service is defined as an interface, with a corresponding implementation encapsulating the business logic. This approach promotes loose coupling and testability, allowing for easy substitution or extension of service implementations as requirements evolve.

C. User Interface and Analytics

The live tracking dashboard in RelengDesk provides users with a comprehensive, real-time overview of ongoing and recent builds across multiple products and platforms.



International Journal of Information Technology (IJIT) – Volume 11 Issue 3, May - Jun 2025

Fig. 4 RelengDesk Live Tracking Dashboard showing real-time build status across multiple products. Product names and other sensitive information have been anonymized for confidentiality.



Fig. 5 Detailed build view and pipeline visualization in RelengDesk.



Fig. 6 Build Time Graph in RelengDesk analytics dashboard.

Note: The product names and certain details in the dashboard screenshot have been modified to protect proprietary information.



Fig. 7 Stage Status Graph in RelengDesk analytics dashboard. For confidentiality reasons, the product names and some metadata in the dashboard screenshots have been anonymized.

As shown in Fig. 4, the dashboard displays a grid of build cards, each representing a product build with key metadata such as build version, trigger time, and platform. Color-coded status indicators and customizable filters allow users to quickly assess the state of each build, identify successful and failed jobs, and focus on relevant products or groups. Users can drill down into individual builds to access detailed information, including job names, build numbers, start and end times, and associated metadata. The build pipeline visualization (see Fig. 5) presents each stage's status and duration, enabling users to track progress, pinpoint bottlenecks, and analyze failures within the build process. This intuitive interface supports rapid feedback cycles, efficient monitoring, and data-driven decision-making for engineering teams.

The analytics dashboard in RelengDesk offers comprehensive visualizations and insights into build performance and trends [2]. As illustrated in Fig. 6, users can view build time graphs that display the duration of recent builds for a selected product, enabling quick identification of outliers and performance bottlenecks. Interactive filters allow users to refine analytics by product, version, branch, platform, and build type, supporting targeted analysis. Additionally, the stage status graph (Fig. 7) breaks down the duration of individual pipeline stages within a specific build, helping teams pinpoint stages that contribute most to overall build time. These analytics features empower engineering teams to monitor trends, optimize build configurations, and make datadriven decisions to improve efficiency and reliability in the release engineering process.

D. Integration and Extensibility

A key innovation in RelengDesk is its centralized access control mechanism, which provides unified authentication and authorization across multiple CI/CD platforms, eliminating the need for individual permissions management. The system also supports customizable views per product, parallel build monitoring, and selective build display, addressing limitations found in traditional plugin-based solutions. The modular and extensible design allows for the integration of additional analytics modules, alternative data stores, or new event processing channels with minimal changes to the existing codebase.

RelengDesk is designed to be platform-agnostic and easily integrable with a variety of external CI/CD systems. A key aspect of this integration is the custom Build Service, which consists of a cluster of Python scripts utilized during the build process. These scripts are embedded within the build pipelines of external systems such as Jenkins, CircleCI, or other CI/CD tools. During the execution of a build, these scripts are triggered at both the stage and job levels. They are responsible for collecting relevant build data and pushing structured event notifications to the RelengDesk backend. This is achieved through API calls or by publishing events to Kafka, depending on the configured channel. The event payloads include detailed metadata such as build identifiers, status, timestamps, platform information, and any custom attributes required for analytics or monitoring. This approach enables seamless integration with heterogeneous build environments, allowing RelengDesk to aggregate and monitor builds across multiple platforms without being tightly coupled to any specific CI/CD tool. The modularity of the Build Service scripts also allows for easy extension or adaptation to new build systems as organizational requirements evolve.

By leveraging lightweight, script-based integration, RelengDesk ensures minimal overhead on the build process while providing comprehensive, real-time visibility into build activities across the enterprise.

V. RESULT & DISCUSSION

A. System Performance and User Experience

The deployment of RelengDesk in enterprise environments has resulted in notable improvements in the efficiency and transparency of release engineering processes. The live tracking dashboard provides real-time visibility into the status of multiple builds and products, enabling engineering teams to monitor progress, promptly identify issues, and respond proactively to failures or bottlenecks. The user interface, featuring customizable views and color-coded status indicators, has been positively received, facilitating rapid feedback cycles and reducing the cognitive load associated with monitoring complex build pipelines.

The analytics dashboard further enhances decision-making by presenting aggregated build statistics, error trends, and resource utilization metrics. This functionality enables teams to identify recurring issues, optimize build configurations, and allocate resources more effectively. The ability to filter and analyze historical data supports continuous improvement and fosters a data-driven culture within release engineering teams [17].

B. Enterprise Integration and Comparative Analysis

RelengDesk's platform-agnostic architecture and scriptbased integration approach have proven effective in heterogeneous enterprise environments [14]. By decoupling build event collection from any specific CI/CD tool, the system supports seamless aggregation and monitoring of builds across multiple platforms. The centralized access control mechanism simplifies permissions management and enhances security, while the modular design allows for straightforward extension and adaptation to evolving organizational requirements.

In comparison to traditional plugin-based monitoring tools and general-purpose analytics platforms, RelengDesk offers several distinct advantages. The unified dashboard consolidates build information from disparate sources, eliminating the need for users to navigate multiple interfaces or manage scattered plugins. The system's support for parallel build monitoring, selective build display, and real-time analytics addresses key limitations of existing solutions, particularly in large-scale, multi-team environments [3].

The adoption of RelengDesk has led to measurable improvements in operational efficiency, including faster issue detection, reduced build verification time, and enhanced collaboration between development and release engineering teams [16]. The system's extensibility ensures that new analytics modules, data sources, or integration points can be incorporated with minimal disruption, supporting ongoing process optimization and alignment with best practices in DevOps and continuous delivery.

C. Limitations and Future Directions

While RelengDesk provides comprehensive build tracking and basic analytics, the system currently lacks advanced predictive capabilities and real-time anomaly detection. The existing analytics implementation focuses primarily on historical data aggregation and basic error counting, without the ability to predict potential build failures or detect unusual patterns in real-time. Future work should focus on implementing machine learning models for predictive analytics and developing real-time anomaly detection algorithms to enhance the system's proactive monitoring capabilities [17].

VI. CONCLUSIONS

This paper has presented the design and implementation of RelengDesk, an enterprise-grade release engineering monitoring and analytics system. By leveraging a modular microservices architecture, event-driven processing, and a platform-agnostic integration approach, RelengDesk addresses key challenges faced in modern release engineering, including the need for real-time build tracking, comprehensive analytics, and seamless integration with heterogeneous CI/CD environments.

The system's live tracking and analytics dashboards provide engineering teams with real-time visibility and actionable insights, supporting rapid feedback cycles and data-driven decision-making. The centralized access control mechanism and customizable user interface further enhance usability, security, and operational efficiency. RelengDesk's extensible design allows for straightforward adaptation to evolving organizational requirements and integration with additional tools and platforms.

Compared to traditional plugin-based and general-purpose monitoring solutions, RelengDesk offers significant advantages in terms of flexibility, scalability, and user experience. Its ability to aggregate build data across multiple platforms, support parallel build monitoring, and provide selective build display addresses critical limitations of existing tools.

In summary, RelengDesk represents a robust, extensible, and high-performance solution for enterprise release engineering, empowering organizations to accelerate software delivery, reduce operational risk, and make informed, datadriven decisions. The system's successful deployment in production environments demonstrates its practical value and effectiveness in real-world scenarios. As organizations continue to adopt cloud-native architectures and microservices, RelengDesk's platform-agnostic approach and scalable design position it as a valuable tool for modern software development teams.

ACKNOWLEDGMENT

The authors would like to express their sincere gratitude to their mentors and colleagues at Adobe Systems for their continuous support and valuable feedback throughout the course of this research. Special thanks are extended to the software engineering team for their assistance with system implementation and testing. The authors also appreciate the encouragement and guidance provided by their families and friends. Finally, the open-source community and the developers of the tools and frameworks utilized in this work are gratefully acknowledged for their contributions, which made this research possible.

REFERENCES

- [1] J. Anderson and S. Park, "Next-Generation Event-Driven Architectures: Patterns and Implementation Strategies," ACM Comput. Surv., vol. 56, no. 2, pp. 1-35, Feb. 2024.
- [2] L. Martinez and R. Chen, "Real-time Analytics in Cloud-Native Applications: A Comprehensive Study," IEEE Trans. Cloud Comput., vol. 12, no. 2, pp. 112-128, Mar. 2024.
- [3] T. Brown and M. Davis, "Event-Driven Systems in the Age of Cloud Computing: Challenges and Solutions," J. Cloud Comput., vol. 13, no. 1, pp. 45-62, Jan. 2024.
- [4] M. Wilson and P. Taylor, "Error Handling in Distributed Systems: A Modern Approach," J. Distrib. Comput., vol. 36, no. 2, pp. 78-92, Apr. 2023.
- [5] A. Johnson and B. Smith, "Build Systems in Modern Software Development: A Comprehensive Analysis," IEEE Trans. Softw. Eng., vol. 49, no. 3, pp. 156-170, Mar. 2023.
- [6] J. Lee and H. Kim, "Data Synchronization in Cloud-Native Distributed Systems: A Contemporary Analysis," IEEE Trans. Cloud Comput., vol. 12, no. 2, pp. 45-62, Feb. 2024.
- [7] S. Hassan, A. Hindle, and E. Stroulia, "Release Engineering in the Era of Continuous Delivery:

Challenges and Opportunities," IEEE Software, vol. 38, no. 2, pp. 54-61, Mar.-Apr. 2021.

- [8] M. Shahin, M. Ali Babar, and L. Zhu, "Continuous Integration, Delivery and Deployment: A Systematic Review on Approaches, Tools, Challenges and Practices," IEEE Access, vol. 8, pp. 39001-39035, 2020.
- [9] J. Soldani, D. A. Tamburri, and W. van den Heuvel, "The pains and gains of microservices: A Systematic grey literature review," Journal of Systems and Software, vol. 146, pp. 215-232, Dec. 2018.
- [10] S. Newman, Building Microservices: Designing Fine-Grained Systems, 2nd ed. Sebastopol, CA, USA: O'Reilly Media, 2021.
- [11] P. Webb, "Spring Boot Reference Documentation," Pivotal Software, 2023. [Online]. Available: https://docs.spring.io/springboot/docs/current/reference/htmlsingle/
- [12] M. Rodriguez and A. Kumar, "Advanced Event Processing with Apache Kafka: A Study of Modern Real-time Data Architectures," IEEE Trans. Cloud Comput., vol. 12, no. 1, pp. 78-95, Jan. 2024.
- [13] A. Gokhale, D. C. Schmidt, and B. Natarajan, "Event-Driven Architecture for Cloud-Native Applications: A Case Study," IEEE Cloud Computing, vol. 8, no. 1, pp. 36-45, Jan.-Feb. 2021.
- [14] D. Taibi, V. Lenarduzzi, and C. Pahl, "Architectural Patterns for Microservices: A Systematic Mapping Study," IEEE Trans. Softw. Eng., vol. 47, no. 11, pp. 2446-2466, Nov. 2021.
- [15] S. Bass, I. Weber, and L. Zhu, DevOps: A Software Architect's Perspective, 2nd ed. Boston, MA, USA: Addison-Wesley, 2022.
- [16] S. Krusche and L. Alperowitz, "Introduction of Continuous Integration and Continuous Delivery in a Legacy Software Project," Empirical Software Engineering, vol. 25, pp. 492-526, 2020.
- [17] S. Amershi et al., "Software Engineering for Machine Learning: A Case Study," in Proc. 41st Int. Conf. Softw. Eng. (ICSE-SEIP), Montreal, QC, Canada, 2019, pp. 291-300, doi: 10.1109/ICSE-SEIP.2019.00042.