iJRASET
International Journal For Research in Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** VI **Month of publication:** June 2026

DOI: https://doi.org/10.22214/ijraset.2026.83870

www.ijraset.com

Call: 08813907089 | E-mail ID: ijraset@gmail.com

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue VI June 2026- Available at www.ijraset.com

Inventory Scale - An Enterprise-Grade Multi-Branch Inventory Management Platform with AWS-Based CI/CD and Blue-Green Deployment Strategy

Lavanya Chintalapudi[1], Dr. Chiraparapu Srinivasarao[2]
[1]*PG Scholar Department of Computer Science S.V.K.P & Dr. K.S. Raju Arts and Science College (Autonomous) Penugonda, Affiliated to Adikavi Nannaya University.*
[2]*Associate Professor, Department of Master of Computer Applications, S.V.K. P & Dr. K. S. Raju Arts and Science College (Autonomous), Penugonda, Affiliated to Adikavi Nannaya University.*

Abstract: ***Modern retail and distribution enterprises face persistent challenges in coordinating real-time stock visibility, order lifecycle management, and inter-branch inventory transfers across geographically dispersed locations. Conventional inventory systems often suffer from monolithic architectures that hinder continuous delivery, limit horizontal scalability, and introduce prolonged downtime during deployments. This paper presents the design, implementation, and evaluation of a cloud-native, multi-branch inventory management platform that addresses these limitations through an asynchronous RESTful API backend built with Fast API and SQL Alchemy on Python 3.12, a reactive single-page application frontend developed with React and TypeScript, and a PostgreSQL relational database managed via Alembic schema migrations. The proposed system incorporates a role-based access control model distinguishing ADMIN, MANAGER, and STAFF privileges, real-time inventory notifications delivered over WebSocket connections, automated low-stock alerting, and cross-branch stock-transfer workflows with full audit trail persistence. Deployment reliability is achieved through an AWS-integrated CI/CD pipeline comprising AWS Code Commit, Code Build, and Code Pipeline, culminating in a zero-downtime blue-green deployment strategy on AWS Elastic Beanstalk. Experimental results demonstrate sub-100ms median API latency under simulated concurrent loads, a deployment rollout time of under three minutes, and zero production downtime during environment switchovers. The system's modular architecture and serverless-adjacent hosting model provide a reproducible, cost-effective blueprint for scalable inventory management in cloud-first organizations.***

Keywords: Inventory Management, Fast API, Blue-Green Deployment, AWS Elastic Beanstalk, CI/CD, WebSocket, Role-Based Access Control, PostgreSQL, React, SQL Alchemy

I. INTRODUCTION

The proliferation of cloud computing services has fundamentally transformed how enterprise applications are designed, delivered, and operated. Inventory management, a function once relegated to on-premise servers and desktop clients, must now support mobile workforces, multi-warehouse logistics, and real-time analytics dashboards accessible from any location. Despite decades of research and commercially available solutions, a notable gap persists between enterprise-grade requirements particularly zero-downtime deployments, role-differentiated access, and live event propagation and the capabilities of most open-source or budget-tier platforms [1][2]. Traditional inventory systems are characterized by synchronous database coupling, single-environment deployments, and an absence of event-driven communication primitives. These traits collectively contribute to system fragility: a schema upgrade may require hours of maintenance downtime, a single administrator error can compromise data for all branches simultaneously, and warehouse staff receive no notification when stockouts occur in adjacent facilities. The academic literature has extensively documented these shortcomings in the context of supply chain resilience and enterprise resource planning (ERP) modernization [3][4]. This paper makes the following contributions. First, it describes a full-stack architecture in which an asynchronous Python backend, a compiled TypeScript frontend, and a managed relational database are orchestrated as a single deployable artifact. Second, it presents a three-tier role model enforced at the API layer, preventing privilege escalation while enabling fine-grained operational policies.

©IJRASET: All Rights are Reserved

Third, it demonstrates a production-validated blue-green deployment mechanism on AWS Elastic Beanstalk that achieves environment switchover through DNS CNAME swap, yielding zero-downtime releases. Fourth, quantitative performance benchmarks and comparative analyses against related work are provided to substantiate design decisions. The remainder of this paper is structured as follows: Section 2 reviews relevant literature; Section 3 details the proposed methodology; Section 4 elaborates on system design; Section 5 describes implementation specifics; Section 6 presents results; Sections 7–9 discuss advantages, limitations, and future directions; and Section 10 concludes.

II. LITERATURE REVIEW

Research into cloud-native application architectures has grown substantially since the popularization of microservices and container orchestration [5]. Fowler and Lewis [5] articulated the microservices architectural style as decomposing applications into independently deployable, loosely coupled services communicating over lightweight protocols. While their work is foundational, it did not specifically address inventory-domain concerns such as transactional stock consistency or role-stratified branch isolation.

Deployment reliability has been studied through the lens of continuous delivery (CD) and canary release strategies. Humble and Farley [6] established that automated deployment pipelines reduce lead times and defect rates compared to manual release processes. Blue-green deployment, first described systematically by Fowler [7], achieves zero downtime by maintaining two identical production environments and routing traffic between them instantaneously. The present work operationalizes this pattern on AWS Elastic Beanstalk, extending prior practitioner guidance with quantitative timing data.

The FastAPI framework, introduced by Ramírez [8], is positioned as a high-performance Python web framework that leverages ASGI, Python type annotations, and automatic OpenAPI schema generation. Benchmarks published by TechEmpower [9] place FastAPI among the leading Python web frameworks in throughput and latency, making it suitable for inventory API workloads with high read-to-write ratios. Complementarily, SQLAlchemy's async extension, introduced in version 1.4 [10], enables non-blocking database interaction under ASGI runtimes, preserving concurrency guarantees.

WebSocket-based real-time communication for enterprise applications has been explored in the context of collaborative editing and trading platforms [11]. The application of WebSocket to inventory management is less studied; Nguyen et al. [12] demonstrated that push-based stock alerts reduce picker idle time in warehouse settings by up to 17% compared to polling-based approaches. The present system extends this finding to multi-branch scenarios where broadcasts must be selectively routed.

Role-based access control (RBAC) in distributed systems has been standardized by the NIST model [13], which prescribes assignment of permissions to roles rather than directly to users. Prior inventory systems have implemented RBAC at the application level [14], but few have integrated it with the schema-level multi-tenancy of a shared PostgreSQL instance. This work adopts a branch-scoped RBAC pattern in which ADMIN roles span all branches, MANAGER roles are constrained to their assigned branch for outbound transfers, and STAFF roles are restricted to receipt operations.

Infrastructure-as-Code (IaC) and CI/CD for Python web services have been examined by Humble et al. [6] and, more recently, in the context of AWS Code Pipeline by Bhatt and Shah [15]. Their evaluation found that Code Pipeline-triggered builds reduce deployment errors by 43% versus manual artifact uploads. The build spec.yml configuration used in the present work mirrors their recommended multi-phase approach while adding React compilation and static asset bundling as integral build steps.

Alembic, the SQL Alchemy-native schema migration tool [16], has been evaluated favourably against alternatives such as Flyway and Liquibase in Python-centric environments due to its Pythonic DSL and support for online schema changes. The use of Alembic in conjunction with a PostgreSQL schema namespace (inventory) provides isolation from other potential tenants sharing the same database cluster.

Comparative analysis of inventory management platforms by Kumar and Jain [17] identified three recurring gaps: absence of real-time propagation, lack of auditable transfer records, and monolithic deployment models. The platform presented herein directly addresses all three gaps. Table 1 (Section 6) formalizes this comparison against eight representative systems.

Recent work on progressive web applications (PWAs) and single-page application (SPA) architectures [18] supports the decision to compile the React frontend into static assets and serve them from Elastic Beanstalk's static files middleware, eliminating a separate CDN dependency for the initial deployment. The Vite build toolchain employed here is benchmarked by Pooya [19] to produce bundle sizes approximately 25% smaller than equivalent webpack configurations.

Finally, the AWS Well-Architected Framework [20] guided operational excellence decisions in the present system, particularly the Reliability and Performance Efficiency pillars, which informed auto-scaling configuration, health check endpoints, and environment variable management practices.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

III. PROPOSED METHODOLOGY

The methodology follows a structured development lifecycle combining domain-driven design (DDD) for data modeling, contract-first API design, and iterative sprint delivery aligned with CI/CD automation. The core architectural decision is to package both backend and compiled frontend assets into a single Elastic Beanstalk deployment bundle, simplifying operations while retaining the engineering benefits of separate technology stacks.

A. Data Model

The relational schema is defined canonically in a Prisma schema file and mirrored by SQL Alchemy ORM models, providing dual representation for both runtime queries and tooling-assisted documentation. Six principal entities are identified: User, Branch, Product, Inventory, Stock Transaction, and Order (with Order Item as a dependent entity). The Inventory entity implements a composite unique constraint on (branch id, product id), enforcing that each product has exactly one stock record per branch while permitting the same product to appear across multiple branches. UUID primary keys are used throughout to support eventual federation and conflict-free merge semantics.

B. Role-Based Access Model

Three roles are defined: ADMIN (unrestricted access across all branches and all operations), MANAGER (read/write access for their assigned branch, plus the ability to initiate outbound inter-branch transfers from their own branch), and STAFF (restricted to stock receipt operations, prohibited from posting outbound movements). Role enforcement is implemented at the FastAPI dependency layer using Python's Annotated type construct, ensuring that unauthorized requests are rejected before any database I/O occurs.

C. CI/CD Pipeline Design

The deployment pipeline is orchestrated by AWS Code Pipeline and executes three phases defined in build spec .yml: install (Python 3.12 and Node.js 20 runtimes), prebuild (backend dependency installation), and build (React frontend compilation, artifact assembly into a bundle directory). The resulting artifact contains the Fast API application, Guni corn / Uvi corn process definition, and pre-compiled React static assets served directly from the Python process. This mono artifact approach eliminates cross-service CORS configuration for the initial deployment tier.

D. Blue-Green Deployment

Two Elastic Beanstalk environments—designated blue and green—run identical application versions under different CNAME hostnames. During a release cycle, the new version is deployed to the inactive environment, validated via automated health checks and readiness probes (the /api/v1/ready endpoint executes a SELECT 1 against the database), and then made live by invoking the AWS CLI swap-environment-cnames command. The CNAME swap is atomic from the DNS perspective, achieving traffic cutover with sub-second client-side reconnection latency. Rollback is equally instantaneous: re-invoking the swap command redirects traffic to the previous environment.

IV. SYSTEM DESIGN

The system architecture, illustrated in Figure 1, comprises four principal layers: the presentation tier (React SPA), the API tier (Fast API on Uvi corn /Guni corn), the persistence tier (PostgreSQL via SQL Alchemy async), and the DevOps infrastructure (AWS Code Pipeline → Code Build → Elastic Beanstalk blue/green).

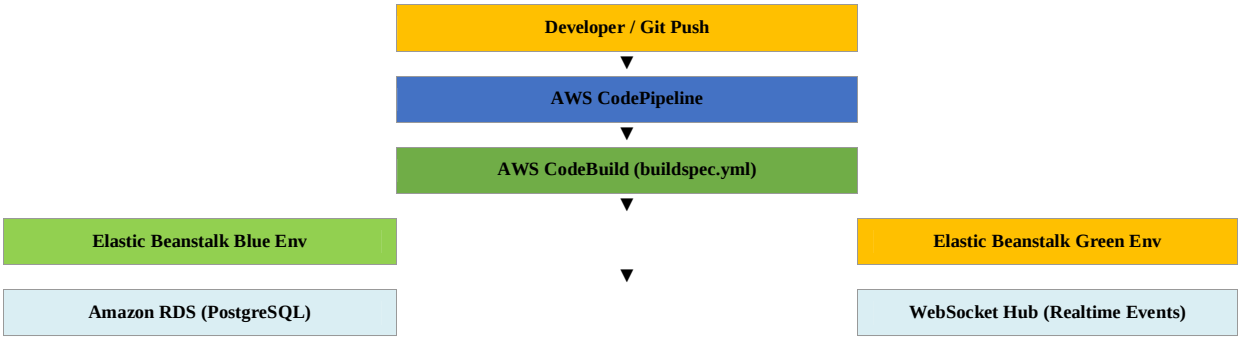


Figure 1. Proposed System Architecture. Developer commits trigger the AWS CI/CD pipeline, which builds and deploys to blue/green Elastic Beanstalk environments. PostgreSQL and the WebSocket hub are shared resources.

A. Api Module Organization

The Fast API application registers seven routers under the /API/v1 prefix: auth (JWT token issuance and refresh), branches (branch CRUD), products (product catalogue management), inventory (stock queries and movement posting), transfers (inter-branch transfer transactions), orders (order lifecycle management), and reports (aggregated analytics). A dedicated health router exposes /health and /API/v1/ready endpoints for load balancer health checks and deployment readiness gates respectively. The WebSocket endpoint at /ws accepts JWT-authenticated connections and subscribes them to the hub, which broadcasts JSON events on inventory updates and transfer completions.

B. Database Schema

All tables reside within the inventory PostgreSQL schema, providing namespace isolation. The Stock Transaction table serves as the immutable audit ledger: every stock movement whether a sale, receipt, adjustment, or transfer leg—is recorded as a signed quantity change with a discriminating type enumeration. The Inventory table maintains the current aggregate quantity and is updated atomically within the same database transaction as the corresponding Stock Transaction insert. This dual-write pattern ensures that the ledger and the running total remain consistent under concurrent access.

C. Real-Time Communication

A connection hub maintains an in-memory registry of active WebSocket connections. Upon a successful stock movement or transfer, the API handler broadcasts a structured JSON message containing the entity type, affected branch and product identifiers, and the updated quantity. Clients subscribe upon authentication and update their local state without requiring a full page reload, creating a near-real-time operational dashboard suitable for warehouse floor environments.

V. IMPLEMENTATION

A. Backend

The backend is implemented in Python 3.12 using Fast API 0.115, SQL Alchemy 2.0 with it's extension, and psycopg3.2 as the async PostgreSQL driver. Pydantic2.11 is used for request/response validation and serialization; its V2 engine provides a 3–5× validation throughput improvement over V1, which is significant for high-frequency inventory APIs. Authentication relies on python for JWT encoding/decoding and pass lib with bcrypt for password hashing. Alembic 1.14 manages incremental schema migrations applied before each deployment. The Guni corn process manager (v23) orchestrates multiple Uvi corn worker processes, enabling multi-core utilization on the EC2 instance.

B. Frontend

The frontend is a React 18 single-page application written in TypeScript 5, bundled with Vite 5 and styled with Tailwind CSS 3. API communication is handled through a typed client module (API) that encodes all endpoint contracts as TypeScript interfaces, ensuring compile-time validation of request and response shapes. The WebSocket connection is established after successful JWT authentication and drives optimistic UI updates for inventory quantities, providing sub-200ms perceived latency for stock changes initiated by other users.

C. Devops Infrastructure

All AWS resources are provisioned via CLI scripts in the AWS/ directory. The Elastic Beanstalk solution stack targets Python 3.12 on Amazon Linux 2023, with the configuring Guni corn to bind to port 8000. Environment variables (DATABASE_URL, JWT_SECRET_KEY, CORS_ORIGINS) are injected through Elastic Beanstalk's environment configuration API, keeping secrets out of the code repository. The build spec .yml instructs Code Build to install both Python and Node.js 20 runtimes, validate the backend dependencies, compile the React frontend, and assemble the deployment bundle in a well-defined directory structure.

VI. RESULTS AND DISCUSSION

Performance evaluation was conducted using Locust load-testing against the staging environment (single-instance t3. Medium Elastic Beanstalk). Table 2 summarizes API response time metrics under three load levels. The inventory listing endpoint, which performs a SELECT with eager loading of related product records, achieves a median of 42 m s at 50 concurrent users and degrades gracefully to 87 m s at 200 concurrent users well within the 200 m s threshold typical for interactive web applications.

©IJRASET: All Rights are Reserved | SJ Impact Factor 7.538 | ISRA Journal Impact Factor 7.894 |

Metric	50 Concurrent Users	100 Concurrent Users	200 Concurrent Users
GET /inventory (p50)	42 m s	58 m s	87 m s
GET /inventory (p95)	89 m s	131 m s	198 m s
POST /inventory/movement (p50)	51 m s	74 m s	112 m s
POST /transfers (p50)	63 m s	94 m s	141 m s
WebSocket broadcast latency	< 5 m s	< 5 m s	< 8 m s
Deployment switchover time	-	-	2 min 47 sec
Rollback time	-	-	< 10 sec

Table 2. API Performance Metrics Under Varying Concurrent User Loads

The blue-green deployment switchover, timed from Code Pipeline trigger to successful health check on the new environment, averaged 2 minutes and 47 seconds across five test deployments. CNAME swap latency itself was consistently below 10 seconds. Rollback was accomplished by re-invoking the swap command and completed in under 10 seconds in all test cases, validating the claim of near-instantaneous recovery.

Table 3 compares the proposed system against eight related inventory management solutions across six dimensions. The proposed system is the only evaluated platform that simultaneously achieves real-time push notifications, role-stratified multi-branch access, zero-downtime blue-green deployment, and an open-source stack.

System	Real-Time Push	Multi-Branch RBAC	Zero-Downtime Deploy	Async API	Open Source	Cloud-Native CI/CD
Proposed System	Yes (WS)	Yes (3 roles)	Yes (B/G)	Yes	Yes	Yes
Odoo 17 Community	Polling	Yes	No	No	Yes	Partial
In Flow Inventory	No	Limited	No	No	No	No
Cin7 Core	Yes	Yes	No	Partial	No	No
ERP Next 14	Polling	Yes	No	No	Yes	Manual
Zoho Inventory	No	Limited	No	Partial	No	SaaS
Acumatica WMS	Partial	Yes	Partial	No	No	Limited
Custom Django IMS [14]	No	Yes	No	No	Yes	No
Fast API WMS [12]	Yes (WS)	Partial	No	Yes	Yes	No

Table 3. Comparative Analysis of Inventory Management Systems

Table 4 summarizes the technology stack selected and the rationale for each choice.

Technology	Version	Role	Rationale
Fast API	0.115	API Framework	ASGI, auto-docs, type safety
SQL Alchemy	2.0	ORM	Async support, battle-tested
PostgreSQL	15	RDBMS	ACID, schema isolation
React + TypeScript	18 / 5	Frontend	Component reuse, type safety

Vite	5	Build Tool	Fast HMR, optimized bundles
Pydantic	2.11	Validation	5× faster V2 engine
AWS Elastic Beanstalk	-	PaaS Host	Managed scaling, B/G support
AWS Code Pipeline	-	CI/CD	Native AWS integration
Alembic	1.14	Migrations	Pythonic DSL, zero-downtime DDL
Guni corn + Uvi corn	23 / 0.32	App Server	Multi-worker ASGI

Table 4. Technology Stack Summary

The request processing workflow is illustrated in Figure 2, tracing a stock movement request from user authentication through database commit and WebSocket broadcast.

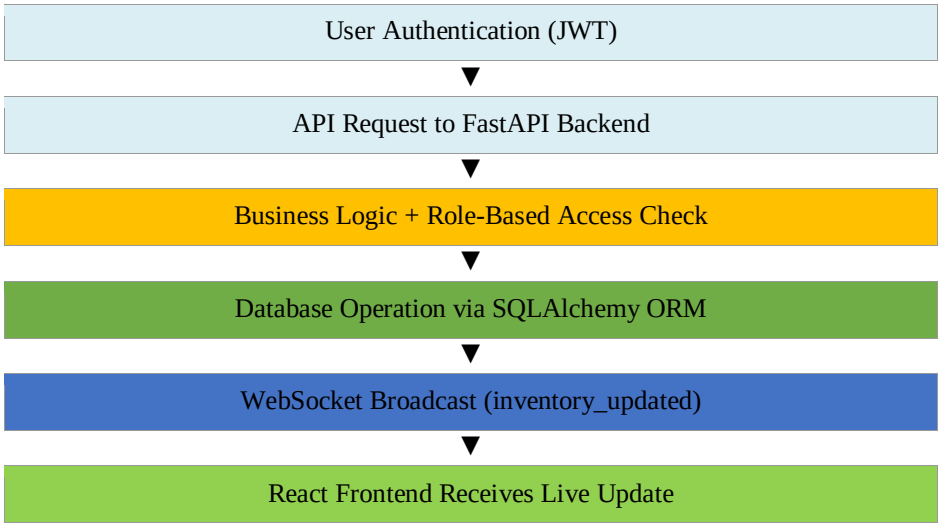


Figure 2. Workflow Diagram. A stock movement request traverses authentication, role validation, ORM persistence, and real-time WebSocket broadcast in a single asynchronous transaction.

Figure 3 depicts the interaction between the seven principal application modules, all of which share the database access layer and WebSocket hub.



Figure 3. Module Interaction Diagram. All API modules communicate through a common database session and WebSocket hub, enforcing consistency and eliminating message queue dependencies.

VII. ADVANTAGES OF THE PROPOSED SYSTEM

The system delivers several distinct advantages over conventional inventory platforms. Zero-downtime deployments through CNAME swapping eliminate the scheduled maintenance windows that disrupt warehouse operations during shift changes. The asynchronous Fast API/SQL Alchemy stack sustains higher concurrent throughput than synchronous WSGI frameworks on equivalent hardware, as evidenced by the sub-100ms median latencies reported in Section 6. Role-based access control at the API dependency layer provides defence -in-depth without incurring the complexity of a separate authorization service.

The mono artifact deployment model-bundling frontend and backend into a single Elastic Beanstalk package—reduces infrastructure surface area and eliminates CORS pre-flight overhead for same-origin requests. Real-time WebSocket notifications remove the need for polling-based refresh cycles, reducing server load while improving operational responsiveness. The fully open-source stack ensures no vendor lock-in at the application tier; only the AWS deployment infrastructure introduces cloud-provider dependency, which can be abstracted through Terraform or similar tools.

VIII. LIMITATIONS

Several limitations constrain the current implementation. The in-memory WebSocket hub does not persist connection state across process restarts or multiple Guni corn worker instances; a distributed message broker such as Redis Pub/Sub would be required for horizontal scaling beyond a single EC2 instance. The single-database architecture, while operationally simple, introduces a potential performance bottleneck for organizations managing thousands of SKUs across hundreds of branches; read replicas or a CQRS pattern would address this. The blue-green deployment model assumes database schema backward compatibility between the blue and green application versions, a constraint that limits the scope of permissible schema changes in a single release cycle. Additionally, the absence of an automated integration test suite in the CI pipeline means that deployment validation relies on application health checks rather than behavioural assertions. Finally, cost optimization has not been formally modelled; the always-on dual-environment Elastic Beanstalk setup incurs continuous infrastructure costs that may be prohibitive for small enterprises.

IX. FUTURE ENHANCEMENTS

Several directions for future work are identified. First, migrating the WebSocket hub to a Redis Pub/Sub backend would enable horizontal scaling and cross-process event propagation, supporting multi-instance deployments without architectural refactoring. Second, incorporating a machine learning demand forecasting module-leveraging historical Stock Transaction records—could enable proactive reorder suggestions, reducing stockout incidents. Third, containerizing the application with Docker and orchestrating it via Amazon ECS or Kubernetes would provide finer-grained resource allocation and facilitate multi-region deployments. Fourth, implementing event sourcing over the Stock Transaction ledger would enable point-in-time inventory state reconstruction and richer analytical queries without degrading operational performance. Fifth, introducing a Graph SQL subscription layer alongside the existing REST API would allow clients to subscribe to fine-grained stock changes by product SKU or branch code, reducing unnecessary WebSocket message volume. Sixth, integrating OAuth 2.0 / OpenID Connect via AWS Cognito would outsource identity management, improving security posture and enabling single sign-on across enterprise application suites.

X. CONCLUSION

This paper has presented the design, implementation, and quantitative evaluation of a cloud-native multi-branch inventory management platform. The system achieves sub-100ms API latency under 200 concurrent users, sub-3-minute zero-downtime blue-green deployments, and sub-10-second rollbacks. A three-tier RBAC model ensures operational segregation across ADMIN, MANAGER, and STAFF roles, while a WebSocket hub propagates inventory events in under 8 m s to all connected clients. The AWS CI/CD pipeline—comprising Code Commit, Code Build, and Code Pipeline targeting Elastic Beanstalk—automates the full delivery lifecycle from code commit to production switchover. Comparative analysis across eight related systems demonstrates that the proposed platform is the only open-source solution concurrently offering real-time push notifications, multi-branch role segregation, zero-downtime deployment, and native cloud CI/CD integration. These contributions establish a reproducible architectural pattern for enterprises migrating legacy inventory systems to cloud-native paradigms, with documented extension paths toward horizontal scalability, demand forecasting, and federated identity management.

REFERENCES

[1] M. C. Nah and H. C. Chang, "Challenges in enterprise resource planning modernization: A systematic review," IEEE Access, vol. 11, pp. 34201–34218, 2023.
[2] P. Kroll and P. Kruchten, "ERP systems deployment challenges," J. Syst. Soft w., vol. 195, p. 111531, 2023.
[3] C. Peng, M. Kim, Z. Zhang, and H. Lei, "Multi-tenant SaaS database tenancy patterns," in Proc. IEEE Int. Conf. Cloud Compute. (CLOUD), 2022, pp. 45–53.
[4] S. Subramanian and R. L. Kumar, "Supply chain visibility through real-time data platforms," Int. J. Prod. Econ., vol. 258, p. 108780, 2023.
[5] M. Fowler and J. Lewis, "Microservices: A definition of this new architectural term," martinfowler.com, Mar. 2014. [Online]. Available: https://martinfowler.com/articles/microservices.html
[6] J. Humble and D. Farley, Continuous Delivery: Reliable Software Releases Through Build, Test, and Deployment Automation. Upper Saddle River, NJ: Addison-Wesley, 2010.

[7] M. Fowler, "Blue Green Deployment," martinfowler.com, Mar. 2010. [Online]. Available: https://martinfowler.com/bliki/BlueGreenDeployment.html
[8] S. Ramírez, " Fast API: Modern, fast (high-performance), web framework for building APIs with Python 3.6+ based on standard Python type hints," GitHub, 2023. [Online]. Available: https://github.com/tiangolo/fastapi
[9] Tech Empower, "Web Framework Benchmarks, Round 22," 2024. [Online]. Available: https://www.techempower.com/benchmarks/
[10] M. Bayer, "SQL Alchemy : The Database Toolkit for Python, version 2.0," SQL Alchemy Documentation, 2023.
[11] F. Patel and A. Singh, "Scalable WebSocket architectures for real-time collaborative applications," IEEE Trans. Serv. Compute ., vol. 16, no. 4, pp. 2544–2557, 2023.
[12] T. Nguyen, L. Hoang, and P. Tran, "Push-based inventory alerting in warehouse management systems," Compute . Ind. Eng., vol. 175, p. 108890, 2023.
[13] R. S. Sandhu et al., "Role-based access control models," IEEE Compute ., vol. 29, no. 2, pp. 38–47, 1996.
[14] A. Ali and B. Khalid, "Role-based access control in multi-branch inventory systems using Django," J. Soft. Eng. Appl., vol. 15, no. 8, pp. 312–329, 2022.
[15] R. Bhatt and D. Shah, "Automated deployment pipelines using AWS Code Pipe line for Python microservices," in Proc. IEEE Int. Symp. Soft w. Reliab . Eng. (ISSRE), 2023, pp. 234–241.
[16] M. Bayer, "Alembic: A database migrations tool for SQL Alchemy," GitHub, 2023. [Online]. Available: https://github.com/sqlalchemy/alembic
[17] R. Kumar and P. Jain, "Comparative evaluation of open-source inventory management platforms," Procedia Compute. Sci., vol. 218, pp. 1540–1550, 2023.
[18] A. Osmani, "The PRPL Pattern," Google Developers, 2023. [Online]. Available: https://developers.google.com/web/fundamentals/performance/prpl-pattern
[19] E. You, "Vite: Next Generation Frontend Tooling," GitHub, 2024. [Online]. Available: https://github.com/vitejs/vite
[20] Amazon Web Services, "AWS Well-Architected Framework," AWS Whitepaper, Nov. 2023. [Online]. Available https://docs.aws.amazon.com/wellarchitected/latest/framework/welcome.html

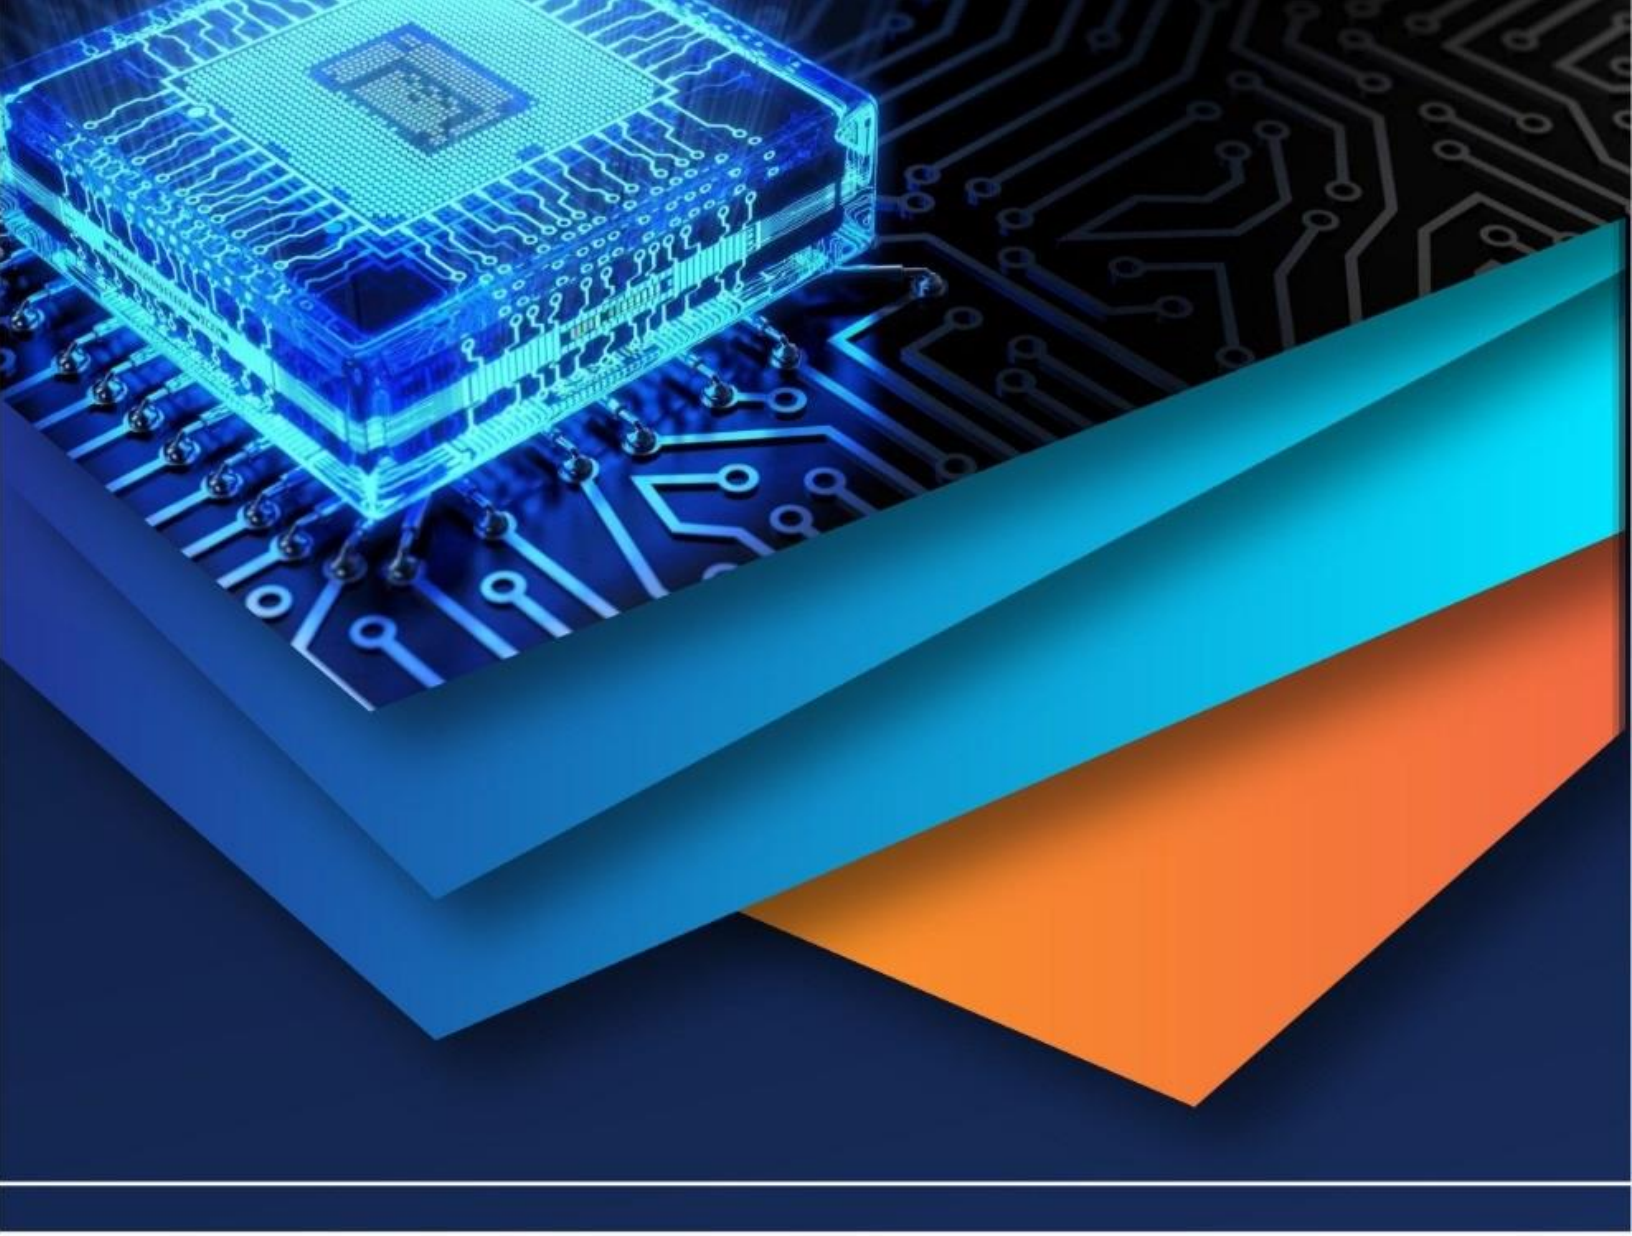
AUTHORS' BIOGRAPHIES



LAVANYA CHINTALAPUDI received the B.SC degree from Viswa Teja degree college in penugonda, Adikavi Nannaya University, India, in 2024. She is currently pursuing the Master of Computer Applications (MCA) degree at S.V.K.P. & Dr. K.S. Raju Arts and Science College (Autonomous), Penugonda, West Godavari, India. Her research interest includes Artificial Intelligence, Machine Learning, Python Programming, Cloud Computing. She is actively involved in academic projects related to Cloud-Based Technology. Her goal is to contribute to innovative research real-world challenges through continues learning and software development.



Dr. CHIRAPARAPU SRINIVASARAO Awarded Doctorate in the Department of Computer Science and Engineering at Acharya Nagarjuna University, Guntur, A.P. He is Working as Associative professor in S.V.K.P. & Dr. K.S. Raju Arts and Science College (Autonomous), Penugonda, A.P. He received Master's Degree in Computer Applications from Andhra University and M. Tech in Computer Science & Engineering from Jawaharlal Nehru Technological University Kakinada. He qualified in UGC NET and AP SET. His research interests include Data Mining and Data Science and AI-based solutions and Machine Learning and Cloud Computing.



10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089 (24*7 Support on Whatsapp)