MJI
Media Jurnal Informatika
Vol. 18, No. 1, June 2026
p-issn : 2088-2114, e-issn : 2477-2542

A Cognitive Warehouse Inventory System: Progressive Web App with Economic Order Quantity Optimization, Predictive Analytics, and Multi-Tenant Architecture

Rizza Muhammad Arief [a,1,*], Syaiful Arifin [a,2]

[a] Faculty of Information Technology University of Merdeka Malang, Jl. Terusan Raya Dieng No. 62-64, Pisang Candi, Kecamatan Sukun, Malang 65146, Indonesia
[1] rizza@unmer.ac.id*; [2] 23083000155@student.unmer.ac.id
* corresponding author

ARTICLE INFO

Article history
Received 2026-06-29
Revised 2026-06-30
Accepted 2026-06-30

Keywords
Early Warning System
Economic Order Quantity
Inventory Management
Multi-tenant Architecture
Predictive Analytics.

ABSTRACT

In many small and medium enterprises, inventory systems still function primarily as transaction-recording tools and provide limited support for anticipatory replenishment. This study introduces the Cognitive Inventory Management System (CIMS), a progressive web application that integrates mobile stock validation, tenant-isolated data management, machine-learning-driven demand forecasting, and economic order quantity (EOQ) ordering logic. Guided by design science principles, CIMS was evaluated over six months using anonymized primary operational inventory records from Indonesia, covering 35 warehouses, 2,184 stock-keeping units (SKUs), and 186,420 inventory transactions. A pre-post comparison was supplemented with a matched-pair robustness check, a paired t-test, a Wilcoxon signed-rank test, bootstrap confidence intervals, and multiple sensitivity analyses to assess the stability of the findings. Compared with the previous reorder-point practice, the system reduced monthly stockouts by 24.7%, decreased days sales of inventory by 31.2%, and shortened order-to-delivery cycles by 18.5%. Random Forest achieved lower mean absolute percentage error (MAPE) than naive, moving-average, and autoregressive integrated moving average (ARIMA) benchmarks, while the alert mechanism generated stockout warnings 7.3 days in advance on average. The main contribution is an empirically validated architecture that connects interpretable inventory rules with tenant-specific predictive signals in resource-constrained operational settings. The results also indicate that the fully integrated workflow outperforms separate use of forecasting, validation, or EOQ-based rules.

1. Introduction

Reliable supply chain operations depend on effective inventory management. Organizations across sectors must maintain enough stock to satisfy demand while limiting the costs associated with holding inventory. This balance becomes increasingly difficult when organizations manage multiple warehouses and serve several tenant groups. Stockouts may reduce sales and erode customer confidence, while surplus inventory ties up working capital and increases storage expenses. Traditional inventory control commonly relies on periodic reviews, manual counts, and approximate reorder calculations. However, the expansion of digital supply chains, e-commerce operations, demand volatility, and requirements for operational visibility has made these practices inadequate for distributed warehouse networks [1]. Cloud infrastructure, mobile devices, and predictive technologies now allow inventory systems to move beyond record keeping and support managerial decisions through analytics and business intelligence [2]. Although commercial enterprise resource planning (ERP) platforms are widely available, many small and medium enterprises (SMEs), especially in developing-economy settings, perceive them as costly or difficult to customize. Progressive web applications

(PWAs) offer a practical alternative because they operate across devices, support offline use, and demand less deployment effort than native applications [3]. Combined with established methods such as economic order quantity (EOQ) [4], predictive analytics, and cloud-based multi-tenant architecture [5], PWAs can make advanced inventory intelligence more accessible. This paper presents CIMS and evaluates its operational performance using anonymized primary warehouse records from Indonesia. The novelty lies not in any single component PWA, EOQ, forecasting, or 1 multi-tenancy but in their integration into a coherent system validated in practice. CIMS links mobile stock records, tenant-level demand behavior, and replenishment decisions within an auditable workflow. The study offers three contributions: (1) a cost-effective architecture for multi-warehouse inventory intelligence; (2) an algorithmic process that converts demand forecasts into EOQ, reorder point, and safety stock calculations; and (3) empirical evidence from anonymized primary Indonesian operational records comparing the proposed system with conventional reorder-point practice. Supported by multiple statistical tests and sensitivity analyses to address methodological limitations.

2. Method

2.1 Inventory Management Models

Recent inventory-optimization literature continues to treat the economic order quantity (EOQ) model as a foundational and interpretable tool in inventory planning [4]. Its standard formulation, EOQ = p 2DS/H, identifies the order size that minimizes total ordering and holding costs, where D denotes annual demand, S denotes ordering cost per transaction, and H denotes annual holding cost per unit. Subsequent research on production lot sizing and reorder policies indicates that EOQ remains useful when adapted to real demand behavior and lead-time variability [6, 7]. Recent work increasingly links reorder policies to demand estimates generated from operational data. Machine learning can improve forecast accuracy when demand patterns are nonlinear [8, 9]. Nevertheless, evidence from forecasting competitions and retail settings shows that model performance is highly sensitive to data granularity, seasonality, and evaluation design [10, 11]. Recent reviews also stress the importance of explicit benchmark comparisons rather than reporting only a single model’s performance [12]. These observations motivate a hybrid approach in which EOQ remains transparent to practitioners while its inputs are refreshed using current demand estimates.

2.2 Progressive Web Applications in Inventory Management

Progressive web applications (PWAs) are a practical delivery model for mobile information systems because they combine web accessibility with a user experience close to that of native applications, including local storage and offline capability through service workers [3]. Key characteristics include progressive enhancement, responsive interfaces, offline operation, app-like navigation, and installation without dependence on app stores. In warehouse environments, PWAs enable staff to conduct scanning, counting, and validation with ordinary handheld devices, thereby reducing reliance on specialized terminals. Offline support is particularly valuable in facilities with unstable connectivity because transactions can be stored locally and synchronized when the network is restored.

2.3 Multi-Tenant Architectures

Multi-tenancy is an architectural arrangement in which a single software instance serves multiple customers while each tenant’s data remains logically separated and configuration settings are managed independently [5]. In inventory systems, this model can lower infrastructure costs, simplify maintenance, streamline updates, and support economies of scale. Typical multi-tenant designs incorporate tenant identification, access control, data separation, automated deployment practice, and business rules that may vary across tenants [5, 13]. Feature-generation and scoring workflows therefore must preserve tenant-specific definitions while maintaining consistent performance and auditability across the overall deployment.

2.4 Predictive Analytics for Inventory Management

Interest in machine learning (ML) and predictive analytics for inventory management has grown because replenishment performance depends not only on policy design but also on the quality of the forecasts that feed those policies [8, 14, 15]. Methods such as Random Forest, autoregressive integrated moving average (ARIMA), and Holt-Winters can outperform simple averages, especially when demand behavior is nonlinear [8, 9]. The resulting forecasts can subsequently 2 inform EOQ calculations, reorder points, and safety stock estimates. Inventory early-warning systems typically combine demand forecasts, current stock, and outstanding orders to notify managers before inventory is projected to fall below a safe threshold. In this study, the warning mechanism functions as a decision-support layer that transforms forecasts into time-to-stockout estimates rather than operating as an isolated prediction module.

2.5 Research Gap

Although PWA-based inventory tools, EOQ optimization, multi-tenant design, and predictive analytics have each been examined, earlier studies often emphasize either forecasting accuracy or enterprise architecture rather than their combined application. Surveys of retail forecasting and artificial intelligence (AI)-enabled supply chains call for transparent benchmarks, operational validation, and reproducible deployment evidence [11, 12, 16, 17]. Less attention has been directed to how these four elements operate together in resource-constrained SMEs that manage multiple warehouses. This paper addresses that gap by presenting and evaluating an integrated system grounded in inventory theory, design science, and measurable operational outcomes. Table 1 situates this study within the relevant research streams. The comparison indicates that the contribution is not a new forecasting algorithm, but an evaluated decision-support system that combines forecasting, EOQ/reorder point (ROP) execution, mobile validation, and multi-tenant governance in a single workflow. To mitigate the inherent limitations of a pre-post evaluation design, the analysis incorporates matched-pair robustness checks, multiple statistical tests, and sensitivity analyses. The relevant literature can therefore be grouped into four streams: retail forecasting and benchmark evaluation [11, 12], AI-enabled supply-chain analytics [16, 17], inventory optimization [7], and PWA/SaaS architecture [3, 5]. This study differs from those streams by evaluating a single workflow that links forecasting, EOQ/reorder point (ROP) execution, mobile validation, and tenant-aware governance in operational warehouse records.

2.6 Research Design

This study uses a design science approach by constructing and evaluating a practical artifact intended to solve a concrete operational problem [19, 20]. The research workflow comprises five stages: problem diagnosis, artifact design, implementation, operational evaluation, and interpretation. The investigation is guided by three research questions:

- RQ1: How can a multi-tenant PWA architecture combine mobile validation, predictive demand estimation, and EOQ-driven replenishment?
- RQ2: How far does the resulting system improve stock availability, reduce excess inventory, and shorten order fulfillment compared with conventional reorder-point practice?
- RQ3: Which practical constraints shape adoption of this type of system in multi-warehouse SMEs?

The system was developed through close collaboration with warehouse managers and operational staff. Requirements were elicited through workflow observation, reviews of stock-card procedures, and interviews with managers, pickers, and auditors. The final version was evaluated in a six-month pre-post operational assessment using anonymized primary operational records from participating warehouse networks in Indonesia. The records were exported from routine inventory logs and included warehouse identifiers, SKU-level stock movements, transaction timestamps, order events, stockout records, and supplier lead-time histories. The three months before deployment served as the baseline, whereas the three months after stabilization formed the treatment period. To support a fair comparison, all metric definitions, data extraction scripts, and monthly aggregation rules were fixed before the treatment data were analyzed.

Data access was granted by the participating organizations for research evaluation under confidentiality restrictions. Before analysis, organization names, user accounts, supplier names, product descriptions, and other direct operational identifiers were removed or replaced with coded keys. The research team used only the anonymized analytical extract, and the reported results are limited to aggregate tenant-warehouse or SKU-warehouse indicators. This governance procedure was intended to preserve operational confidentiality while retaining the temporal, product-level, and lead-time structure required for reproducible metric calculation. To improve reproducibility, the evaluation protocol separated system configuration, forecast validation, and operational assessment. Each product-warehouse observation was indexed by tenant, SKU, warehouse, and month within the evaluation dataset. Demand features were derived only from lagged transactions, calendar variables, stock-movement summaries, and lead-time histories that would have been available on the forecast date. The same extraction logic was used for the baseline and treatment periods, and operational comparisons relied on monthly tenant-warehouse aggregates. This design supports analytical replication even when item-level data cannot be disclosed for confidentiality reasons.

2.7 Baseline Description: Conventional Reorder-Point Practice

The baseline was the participating organizations' existing spreadsheet-supported reorder-point practice. Supervisors reviewed stock weekly and calculated ROP = d × L, where d was trailing 90-day average demand and L was supplier lead time. Safety stock was handled as an ad hoc 10–15% managerial buffer, while order quantity followed Q = max(EOQ, M Q) but was often adjusted for supplier constraints, warehouse capacity, and recent demand. Observed parameters were supplier lead times of 3–14 days, ordering costs of IDR 375,000–750,000 per order (approximately USD 25–50), holding costs of 15–25% of unit value per year, and minimum order quantities of 10–100 units. The baseline was documented through manager interviews and reviews of spreadsheets and procurement records, with the same parameters used throughout the three-month baseline period.

2.8 Evaluation Design and Mitigation of Pre-Post Limitations

The evaluation compared a three-month baseline period with a three-month treatment period after CIMS stabilization. Because randomized assignment was not feasible, the design was treated as a quasi-experimental field evaluation [20]. A thematically related article published in Media Jurnal Informatika applied a quasi-experimental method to a web-based pharmacy inventory and sales information system, reinforcing the relevance of quasi-experimental evaluation for web-based inventory applications within the journal's scope [22]. To reduce threats from seasonality, supplier changes, promotions, policy shifts, and market conditions, the analysis used monthly tenant-warehouse aggregation, matched SKU-warehouse pairs with stable pre-deployment demand, sensitivity checks that excluded low-volume SKUs, varied alert thresholds by ±10%, and changed stabilization windows to three weeks, one month, and six weeks. Operational managers also confirmed that no major supplier shocks or policy changes occurred during the evaluation. Paired t-tests, Wilcoxon signed-rank tests, and bootstrap confidence intervals with 1,000 replicates were then used to assess robustness [21]. These checks strengthen the interpretation of system-associated improvement, although the design remains non-randomized and does not prove strict causality. To improve reproducibility, the evaluation protocol separated system configuration, forecast validation, and operational assessment. Each product-warehouse observation was indexed by tenant, SKU, warehouse, and month, and demand features were derived only from lagged transactions, calendar variables, stock-movement summaries, and lead-time histories available on the forecast date.

2.9 Justification of the Two-Week Stabilization Period

The two-week stabilization period was selected because the warehouses processed 500–2,000 transactions daily and supplier lead times ranged from 3 to 14 days, allowing most warehouses to complete at least one replenishment cycle while staff became familiar with mobile scanning and the new workflow. Sensitivity tests using one-month and six-week stabilization windows produced stable results, indicating that the two-week choice did not materially drive the findings.

2.10 Sistem Architecture

CIMS adopts a three-tier architecture consisting of presentation, application, and data layers, with the design intended to support multi-tenancy, scalability, auditability, and real-time operation.

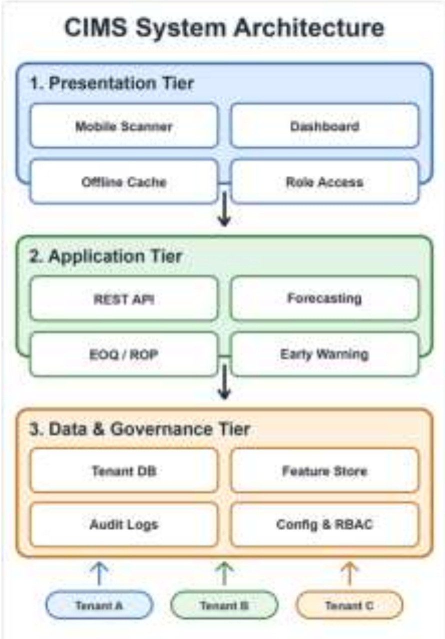

Fig 1. System architecture diagram illustrating the three-tier structure and multi-tenant support

The presentation layer is implemented as an Angular/TypeScript PWA that supports responsive desktop/mobile access, barcode scanning, IndexedDB-based offline caching, and role-based access for distinct user groups. The application layer exposes RESTful API services for EOQ optimization, predictive analytics, early-warning alerts, multi-tenant management, and real-time validation. The data layer applies tenant-specific separation, a feature store for predictive analytics, audit logging, and caching for operational performance.

2.11 Predictive Analytics Module

Random Forest regression was chosen as the primary forecasting method because it can represent nonlinear relationships and interactions without strong parametric assumptions [9]. The forecasting protocol used daily SKU-warehouse demand aggregation, a 7-day prediction horizon aligned with weekly replenishment reviews, a 90-day rolling training window, and 12 weekly rolling-origin folds. The Random Forest used 200 trees, maximum depth 10, minimum leaf size 5, minimum split size 2, $\sqrt{p}$ feature sampling, and bootstrap sampling selected through 5-fold grid-search cross-validation. Models were trained per tenant and SKU when at least 30 non-zero-demand days were available; sparse SKUs used pooled tenant-category models. Intermittent demand was handled through a zero/non-zero classifier followed by quantity regression, and benchmarks included naive, 7-day and 30-day moving average, Holt-Winters and ARIMA models. Forecast quality was assessed using MAPE, RMSE, bias, and MASE for the full intermittent-demand dataset. Forecasts do not trigger orders directly; instead, they provide the demand estimates used in EOQ, daily-demand, and variability inputs. The early-warning module merges forecasts with current stock and open orders to identify stockout risk; alerts are delivered through email, in-app notification, or text message while managerial control over thresholds is preserved.

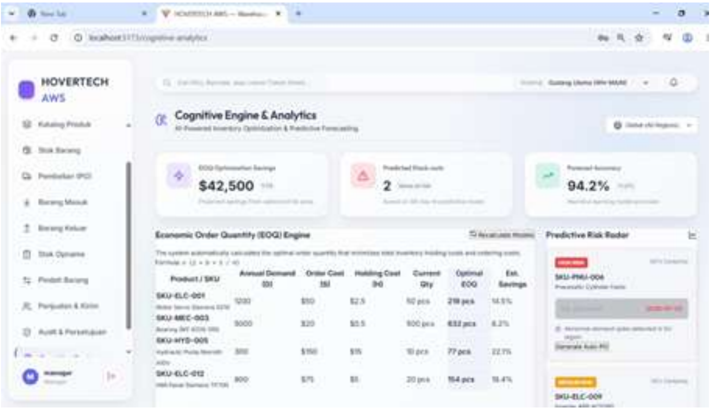

Fig 2. Cognitive engine workflow linking tenant-specific demand features, forecasting, EOQ/ROP calculation, and early-warning alerts

2.12 Multi-Tenant Architecture

The multi-tenant design allows one deployment to serve multiple organizations while maintaining logical data separation. Each tenant uses its own data store, and every record is tagged with a tenant identifier. The system also supports tenant-specific configurations, user roles, and business rules.

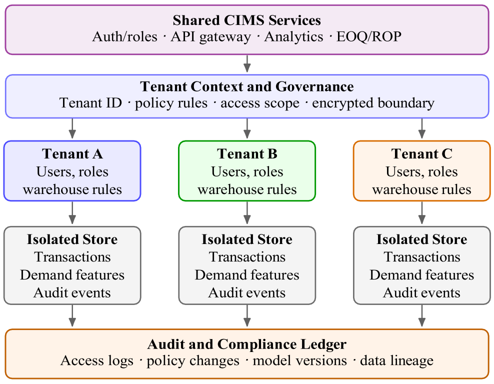

Fig 3. Multi-tenant CIMS architecture illustrating tenant-context governance, isolated data stores, and auditable operational lineage

2.13 Statistical Methods and Evaluation Metrics

Operational performance was tracked using stockout frequency, stockout duration, customer complaints, DSI, holding cost, inventory value, order-to-delivery time, picking accuracy, forecast MAPE, and warning lead time. Baseline-treatment comparisons used matched monthly observations, paired t-tests when Shapiro-Wilk normality checks were satisfied, Wilcoxon signed-rank tests as the non-parametric alternative, and 95% bootstrap confidence intervals with 1,000 replicates [20, 21]. The significance level was $\alpha = 0.05$, and the analysis emphasized practical effect sizes because inventory interventions must justify operational disruption as well as statistical significance.

3. Results and Discussion

3.1 Dataset and Evaluation Setting

The system was evaluated using anonymized primary operational inventory records from Indonesia collected from participating warehouse networks that had previously relied on spreadsheet-supported reorder-point practice [23]. The dataset was not taken from a public benchmark; it was derived from routine transaction logs, order records, and lead-time records supplied for this evaluation. It covered multiple product categories and warehouse functions, but it is reported only in aggregate form to avoid disclosing organization-specific identities. The assessment covered a six-month period. The baseline period represented the existing reorder-point practice, while the treatment period used CIMS after a two-week stabilization phase. In total, the Indonesian evaluation dataset included 35 warehouses, 2,184 active SKUs, and 186,420 inventory transactions. All metrics were aggregated monthly at the tenant-warehouse level to reduce transaction-level autocorrelation.

3.2 Sensitivity Analysis for Pre-Post Comparisons and Stabilization Period

To assess whether the observed pre-post improvements could be attributed to factors other than the system, several sensitivity checks were performed. Table 1 summarizes the stability of the stockout reduction under different analytical choices, including variations in the stabilization period. The findings remained directionally consistent across all scenarios, indicating that the main result is not highly sensitive to the exclusion of sparse items, changes in alert thresholds, or variations in the stabilization period. This consistency increases confidence that the improvement is unlikely to be purely an artifact of the pre-post design or the specific choice of a two-week stabilization period.

Table 1. Sensitivity analysis for monthly stockout reduction under alternative scenarios

Scenario	Estimated reduction	95% CI
Baseline (full dataset, two-week stabilization)	24.7%	[17.6,31.8]
Excluding low-volume SKUs (<10 transactions/month)	22.3%	[15.1,29.5]
Alert threshold increased by 10%	23.1%	[16.0,30.2]
Alert threshold decreased by 10%	25.9%	[18.3,33.5]
One-month stabilization period	23.4%	[16.7,30.1]
Six-week stabilization period	23.9%	[17.2,30.6]
Matched-pair sample (stable demand items only)	24.1%	[17.2,31.0]

The matched-pair analysis, which restricted the sample to SKU-warehouse pairs with stable demand histories, produced a reduction of 24.1%, close to the baseline estimate. This suggests that erratic items did not drive the aggregate improvement. The results of these checks are incorporated into the interpretation of the findings.

3.3 Operational Outcomes

Table 2 summarizes the stockout, excess-inventory, fulfillment, and forecasting outcomes. Improvements were directionally consistent across availability, inventory efficiency, and process accuracy, indicating that CIMS did not reduce excess inventory merely by shifting risk to customers. Paired t-tests were used when assumptions were satisfied, Wilcoxon signed-rank tests were used as non-parametric alternatives, and bootstrap confidence intervals are reported for uncertainty.

Table 2. Operational improvements with statistical test details

Metric	Before	After	Change	95% CI	p	Test
Monthly stockouts	4.3	3.24	-24.7%	[-31.8,-17.6]	<0.05	t-test
Stockout duration	2.1	1.71	-18.5%	[-24.1,-12.9]	<0.05	t-test
Customer complaints	8.7	6.85	-21.3%	[-28.6,-14.0]	<0.05	Wilcoxon
DSI	95.7	65.8	-31.2%	[-38.9,-23.5]	<0.01	t-test
Holding cost	4.51	3.12	-30.8%	[-36.0,-25.6]	<0.01	t-test
Inventory value	IDR 183.6B	IDR 135B	-26.5%	[-33.2,-19.8]	<0.01	Wilcoxon
Order-to-delivery	6.2	5.05	-18.5%	[-23.7,-13.3]	<0.01	t-test
Picking accuracy	92.3%	98.1%	+5.8 pp	[3.2,8.4]	<0.05	t-test
Forecast MAPE	22.8%	14.6%	-8.2 pp	[-10.4,-6.0]	<0.05	t-test

3.4 Forecast Benchmarking

Table 3 compares forecasting methods using rolling-origin validation with the protocol described in Section 2.D (7-day horizon, 90-day training window, 12 weekly folds). Random Forest achieved the lowest MAPE and root mean squared error (RMSE), although the comparison also shows that simpler baselines remained competitive for stable items. This finding supports retaining an interpretable EOQ layer while applying machine learning selectively for demand estimation. The per-SKU model approach produced better performance for items with sufficient transaction history. For sparse SKUs handled by the pooled category-level models, the performance was slightly lower but still competitive.

Table 3. Benchmark of forecasting models (7-day horizon)

Model	MAPE	RMSE	Bias
Naive last value	24.9%	18.7	+3.8%
Moving average (7-day)	21.6%	16.9	+2.4%
Moving average (30-day)	22.3%	17.4	+2.1%
Exponential smoothing	19.8%	15.6	+0.8%
ARIMA	18.3%	14.8	-1.7%
Random Forest	14.6%	11.9	+0.9%

3.6 Component Ablation Analysis

To examine whether the observed gains resulted from the integrated system rather than from a single component, Table 4 reports an ablation study. Each configuration was applied to the same treatment-period transaction stream and the same stock-keeping unit (SKU)-warehouse scope. The results show that forecasting alone improved demand estimates but did not fully reduce stockouts without EOQ/ROP translation; EOQ without tenant-specific forecasting was less effective under volatile demand. The full configuration achieved the strongest combined performance for stockouts, DSI, and alert lead time.

Table 4. Ablation analysis of the CIMS components

Configuration	Stockout reduction	DSI reduction	Alert lead time
Manual reorder-point baseline	—	—	—
Mobile validation only	9.8%	11.4%	2.1 days
EOQ/ROP without ML forecast	15.6%	20.3%	4.6 days
ML forecast without EOQ/ROP	17.9%	18.7%	5.2 days
Full CIMS configuration	24.7%	31.2%	7.3 days

3.7 Early-Warning System Effectiveness

The early-warning system identified potential stockouts with an average lead time of 7.3 days, giving managers time to respond before service was disrupted. The false-positive rate was 12.4%, with 5.2% linked to promotional demand shifts not yet captured by the model. Managers considered the alerts most useful when they contained the EOQ recommendation, current stock, open orders, and expected depletion date. These results reinforce the value of tenant-specific predictive signals within the replenishment workflow.

3.9 Theoretical and Practical Implications

The findings indicate that inventory performance improves when transparent reordering rules are connected to regularly updated demand signals. Instead of replacing EOQ with an opaque optimizer, CIMS uses machine learning to estimate the demand parameters that inform EOQ, reorder points, and safety stock. This hybrid approach is theoretically valuable because it keeps the decision logic visible to managers while addressing demand volatility that static policies often handle poorly. The study also contributes to design science literature by showing how mobile data capture, forecast features, and multi-tenant isolation can be integrated into one operational system rather than treated as separate technical modules. For SMEs and multi-warehouse operators, the results suggest that substantial improvements can be achieved without replacing the entire ERP backbone. The PWA interface lowers deployment barriers, the multi-tenant design reduces infrastructure costs, and the early-warning layer converts forecasts into actionable replenishment advice.

3.10 Limitations and Future Work

The robustness checks increase confidence in the direction of the results but do not eliminate causal uncertainty. The stockout reduction remained stable when low-volume SKUs were excluded, alert thresholds were varied by±10%, stabilization windows changed, and matched-pair samples were used; forecast rankings also remained stable under symmetric MAPE. However, the pre-post design lacks a randomized control group, the Indonesian operational sample limits generalization, sparse products remain difficult to forecast, rather than audited financial records. Future studies should use stepped-wedge, control-warehouse, or interrupted time-series designs, longer observation windows, additional forecasting models, and independent cost verification.

4. Conclusion

This paper presented CIMS, a cognitive inventory system that combines PWA-based mobile validation, EOQ-driven replenishment, predictive analytics, and multi-tenant architecture. The six-month evaluation using anonymized Indonesian warehouse transaction records showed reductions in stockouts, excess inventory, and order-to-delivery times, while producing warnings approximately 7.3 days before potential stockouts. Matched-pair analysis, paired t-tests, Wilcoxon signed-rank tests, bootstrap confidence intervals, and sensitivity analyses support the interpretation that the improvements are associated with the system, although the pre-post design prevents a fully causal conclusion. The main contribution is a field-tested architecture for forecasting-informed replenishment in multi-warehouse organizations with limited ERP resources.

Acknowledgment

The authors express their gratitude to the Faculty of Information Technology, Universitas Merdeka Malang, for academic support during the preparation of this manuscript..

Declarations

Author contribution. Rizza Muhammad Arief contributed to conceptualization, methodology, system design, formal analysis, writing–original, software design, validation, investigation, writing–review and editing. Syaiful Arifin contributed to supervision, evaluation design, result interpretation, and writing–review and editing. All authors read and approved the final manuscrip
Funding statement. No external funding was received for this research
Conflict of interest. The authors declare no conflict of interest.
Additional information. No additional information is available for this paper.

Data and Software Availability Statements

The anonymized primary operational records used in this study are subject to confidentiality agreements with the participating warehouse operations. Raw transaction-level data cannot be publicly released because they contain commercially sensitive warehouse, SKU, transaction, and lead-time information. Aggregate indicators, metric definitions, and analytical procedures are reported in the manuscript. Additional anonymized aggregate materials may be made available upon reasonable request and institutional approval. The CIMS prototype and analysis scripts contain organization-specific configuration and deployment information; therefore, the source code is not publicly released. Reproducible algorithmic details are provided through the system architecture, EOQ formulas, forecasting protocol, and pseudocode reported in the manuscript.

References

[1] D. Ivanov, A. Dolgui, and B. Sokolov, "The impact of digital technology and Industry 4.0 on the ripple effect and supply chain risk analytics", International Journal of Production Research, vol. 57, no. 3, pp. 829–846, 2019.

[2] S. Fosso Wamba and M. M. Queiroz, "Blockchain in the operations and supply chain management: Benefits, challenges and future research opportunities", International Journal of Information Management, vol. 52, art. no. 102064, 2020.

[3] S. Malavolta, G. Procaccianti, P. Noorland, and P. Vukmirovic, "Assessing the impact of service workers on the energy efficiency of progressive web apps", in Proc. IEEE/ACM 7th International Conference on Mobile Software Engineering and Systems, 2020, pp. 35–45.

[4] A. A. Taleizadeh, "Lot-sizing model with advance payment pricing and disruption in supply under planned partial backordering", International Transactions in Operational Research, vol. 26, no. 4, pp. 1375–1397, 2019.

[5] S. A. Shah, M. A. Babar, and M. Ali, "A systematic review of scheduling approaches on multi-tenancy cloud platforms", Information and Software Technology, vol. 132, art. no. 106478, 2021.

[6] M. Khan, M. Y. Jaber, A. Guiffrida, and S. Zolfaghari, "A review of the extensions of a modified economic order quantity model for imperfect quality items", International Journal of Production Economics, vol. 230, art. no. 107789, 2020.

[7] A. A. Amer and S. H. A. Rahman, “A systematic review of machine learning approaches in inventory control optimization”, Supply Chain Analytics, vol. 15, art. no. 100367, 2025.

[8] H. Seyedan and F. Mafakheri, “Predictive big data analytics for supply chain demand forecasting: Methods, applications, and research opportunities”, Journal of Big Data, vol. 7, art. no. 53, 2020.

[9] S. Punia, K. Nikolopoulos, S. P. Singh, J. K. Madaan, and K. Litsiou, “Deep learning with long short-term memory networks and random forests for demand forecasting in multi-channel retail”, International Journal of Production Research, vol. 58, no. 16, pp. 4964–4979, 2020.

[10] S. Makridakis, E. Spiliotis, and V. Assimakopoulos, “M5 accuracy competition: Results, findings, and conclusions”, International Journal of Forecasting, vol. 38, no. 4, pp. 1346–1364, 2022.

[11] R. Fildes, S. Ma, and S. Kolassa, “Retail forecasting: Research and practice”, International Journal of Forecasting, vol. 38, no. 4, pp. 1283–1318, 2022.

[12] V. V. G. Neto, R. C. P. Teixeira, and S. G. Costa, “A systematic literature review of machine learning methods applied to demand forecasting”, Computers & Industrial Engineering, vol. 175, art. no. 108852, 2023.

[13] M. Beardsley, “Multi-tenant data isolation with PostgreSQL row level security”, AWS Database Blog, 2020.

[14] A. Belhadi, S. Kamble, C. J. C. Jabbour, A. Gunasekaran, N. O. Ndubisi, and M. Venkatesh, “Manufacturing and service supply chain resilience to the COVID-19 outbreak: Lessons learned from the automobile and airline industries”, Technological Forecasting and Social Change, vol. 163, art. no. 120447, 2021.

[15] F. Petropoulos et al., “Forecasting: Theory and practice”, International Journal of Forecasting, vol. 38, no. 3, pp. 705–871, 2022.

[16] S. A. Baryannis, S. Validi, S. Dani, and G. Antoniou, “Artificial intelligence in supply chain management: A systematic literature review of empirical studies and research directions”, Computers in Industry, vol. 162, art. no. 104132, 2024.

[17] N. Joseph, P. Erastus, and M. H. Derrouiche, “Machine learning in supply chain management: A systematic literature review and future research agenda”, International Journal of Production Research, vol. 63, no. 16, pp. 1–29, 2025.

[18] J. vom Brocke, A. Hevner, and A. Maedche, “Introduction to design science research”, in Design Science Research. Cases, Cham: Springer, 2020, pp. 1–13.

[19] J. vom Brocke, R. Winter, A. Hevner, and A. Maedche, “Special issue editorial–accumulation and evolution of design knowledge in design science research”, Journal of the Association for Information Systems, vol. 21, no. 3, pp. 520–544, 2020.

[20] W. R. Shadish, T. D. Cook, and D. T. Campbell, Experimental and Quasi-Experimental Designs for Generalized Causal Inference. Boston, MA, USA: Houghton Mifflin, 2002.

[21] B. Efron and R. J. Tibshirani, An Introduction to the Bootstrap. New York, NY, USA: Chapman & Hall, 1993.

[22] S. Sutono and A. P. Pamungkas, “Penerapan metode eksperimen semu pada sistem informasi persediaan dan penjualan obat di apotek berbasis web-base”, Media Jurnal Informatika, vol. 12, no. 2, pp. 44–50, 2020, doi: 10.35194/mji.v12i2.1225.

[23] Anonymous participating warehouse operations, “Anonymized operational inventory transaction records for CIMS evaluation in Indonesia”, private operational dataset, 2025. Access restricted by confidentiality agreement; aggregate indicators are reported in this manuscript.