Optimizing Retail Application Performance: A Systematic Review of Monitoring Tools, Metrics, And Best Practices

[1] Suresh Gangula
[1] Software Engineer, Nike, Inc., OR, USA

Received: 18th Nov 2025 | Received Revised Version: 28th Nov 2025 | Accepted: 19th Dec 2025 | Published: 02th Jan 2026

Volume 08 Issue 01 2026 | Crossref DOI: 10.37547/tajet/Volume08Issue01-02

Abstract

The increasing popularity of digitization in retail has emphasized application performance as a critical component of customer satisfaction, business continuity, and operational efficiency. Retail applications today operate in hyper-dynamic environments, where performance can be introduced by concerns such as workload peaks, scaling cloud infrastructure performance, spikes in transactions, and real-time data processing. Poor performance erodes user experience, depletes objectives, and lowers competitive benefit. Although many surveys of cloud or performance monitoring exist, they do not incorporate retail contexts that emphasize customer facing KPIs, seasonal demand variability, and other omnichannel complexities. The current paper seeks to address this gap in research by providing a systematic review of 45 peer-reviewed studies (2015–2025) specific to monitoring tools, optimization frameworks, and best practices in retail performance monitoring. The findings in the review synthesize performance monitoring platforms (i.e. Splunk, Datadog, New Relic), log management tools, key performance indicators, performance continuous monitoring, performance visualization, and performance integration into retail systems. The paper also offers comparison evaluations of Amazon and Google Cloud monitoring offerings, as well as limitations for optimization planning, and the effect of intelligent analytics to improve scalability and resiliency. The paper contributes by connecting scholarly viewpoints with practical suggestions, providing a distinctive road map for practitioners and academics to develop and enhance high-performance retail applications in cloud-based corporate settings.

Keywords: Retail Application Performance; Splunk; Datadog; New Relic; Data Visualization, Monitoring.

© **2026** Suresh Gangula. This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0). The authors retain copyright and allow others to share, adapt, or redistribute the work with proper attribution.

Cite This Article: Gangula, S. (2026). Optimizing Retail Application Performance: A Systematic Review of Monitoring Tools, Metrics, And Best Practices. The American Journal of Engineering and Technology, 8(01), 07–19. https://doi.org/10.37547/tajet/Volume08Issue01-02

1. Introduction

The rapid growth of digital technologies has drastically transformed the retail landscape to a point that brands depend on execution at scale that's application-driven from data that ensures a better integrated experience for customers. Today's retail applications are required to take on a very large set of functionalities that consist of inventory visibility, personalized recommendations, dynamic pricing, fraud protection, and omnichannel order management [1]. The optimization of performance has now become a requirement for reliability, scalability, and customer satisfaction as the number of users, transactions, and data streams are increasing [2,3]. Even minor performance issues can lead to revenue loss, a compromised user experience, and declining competitiveness in an already hypercompetitive environment that is always increasing customer expectations for user experience [4]. Retail organizations are challenged in their environment because of the dynamic demand, seasonal spikes that can change, and unpredictable user behavior [4]. Traditional application architectures that are based on monolithic deployments, static resource provisioning, cannot solve those issues effectively. Performance bottlenecks may result from inefficiently executing code, non-optimized database queries, network latency, and cloud-based technology's inability to scale quickly enough [5]. Therefore, optimizing performance for retail applications is a systematic approach that will require more advanced monitoring, intelligent resource provisioning, and adaptive scaling ability. Existing surveys and reviews in the field of performance monitoring primarily address general cloud applications or IT infrastructures, focusing on latency, throughput, and resource utilization. However, they rarely address the retail-specific performance landscape, where unpredictable demand spikes, high transaction volumes, seasonal variability, and omnichannel operations present distinctive challenges. Moreover, prior works seldom provide a comparative evaluation of commercial monitoring platforms (e.g., Splunk, Datadog, New Relic) alongside business-driven KPIs such as conversion rates, cart abandonment, and customer retention. This creates a critical research gap in connecting technical monitoring metrics with retail business outcomes.

Performance monitoring and observability platforms are designed to capture a deeper understanding of an application’s health, transactions, and engagement with users. These performance monitoring applications provide an enterprise a single-point-of-reference to uncover inefficiencies and be more proactive about performance enhancements. Furthermore, performance monitoring solutions that leverage predictive analytics and machine learning allow for better anomaly detection and predicting future performance and uptime. Retail application performance optimization is need that goes beyond mere technical effectiveness; it matters to the overall customer experience because a more responsive retail application encourages more engagement, purchase conversions, and brand loyalty. For instance, recent reviews of application improvements and measurement delays. Even a one second delay in latency for a page load has measurable reductions in customer retention and transactions. It should be noted; performance optimization does improve a technical company's issues within an improved performance context; furthermore, it can be quantified as delivering both business improvements.

While there have been several studies examining isolated elements of retail performance e.g. supply chain KPIs, data analytics adoption, fraud detection frameworks, to date there has not been a unified review that systematically examines the optimization of retail applications performance through monitoring tools, metrics, and best practices. Past reviews, to date, have focused either on supply chain efficiencies of omnichannel retailing, on e-commerce adoption, or integration of big data, without the organization alite modern observability platforms (Splunk, Datadog, New Relic), artificial intelligence analytics, and cloud-native monitoring practices contribute to application resilience and scalability. Past work has also rarely directly compared academic frameworks with commercial and simply pointed out challenges related to scalability, integration, and affordability in an applied context. This creates a double challenge where researchers lack a coherent framework for studying application-level performance, and practitioners lack a literature synthesis around selecting and integrating monitoring strategies across omnichannel retail systems.

To address this challenge, this paper presents a systematic review of 45 peer-reviewed research articles published between 2015 and 2025 and synthesizes methods, monitoring frameworks, performance metrics, visualization practices, optimization strategies, and so on. In contrast to previous disjointed research, this review provides a thorough cross-sectional view of retail application performance, highlights important constraints, and suggests future paths for AI-driven predictive analytics, hybrid monitoring, intelligent optimization frameworks that can improve scalability, and customer experience in digital retail ecosystems. The goal of this paper is to provide a thorough examination of the techniques, tools, and methods for optimizing retail application performance. Accordingly, it will review architectural approaches, monitoring frameworks, performance testing techniques, and optimization practices - all in conjunction with their relevance to large-scale retail systems. The overall goal is to provide information and insight to both practitioners and researchers in seeking ways to improve the efficiency, resilience, and scalability of retail applications to engage with contemporary digital commerce. The structure of the remainder of this paper is as follows. Section 2 reviews performance monitoring tools, including an initial overview (2.1), data collection and log management (2.2), metrics and key performance indicators (2.3), real-time monitoring and alerts (2.4), data visualizations and dashboards (2.5), integration into retail systems (2.6), and a summary comparison of tools (2.7). Section 3 considers the challenges and limitations within performance optimization in retail contexts. Section 4 presents real case studies of performance monitoring and optimization practices. Section 5 discusses the findings in detail and their implications on practice. Finally, Section 6 concludes the paper and suggest potential future research directions.

2. Literature Review

2.1. Overview of Performance Monitoring Tools

In 2016, Harrauer and Schnedlitz [11] examined the relationship between information distribution and interpretation on the retail sales floor, thus providing some new contributions of practice theory to management control and performance measurement in complex settings. Although they take a problem-centered qualitative approach and conduct interviews in two contexts (the U.S. and Europe), the authors acquire information from a total of 22 interviewees with representative coverage from retail areas as well as managerial levels, particularly regarding store management.

In 2023, Parisa et al. [12] explored where Artificial Intelligence (AI)-driven Zero Trust Security (ZTS) models can be imposed in retail cloud ecosystems. Our emphasis is on ZT by proposing an intelligent Zero Trust framework that harnesses AI capabilities for ongoing authentication, anomaly detection, and automated threat response. To evaluate how potent the zero trust solution the authors of ZT are conducting simulations and case studies against conventional cyberspace security measures.

In 2025, Chen et al. [13] offered a new Order Fulfillment Process (OFP) model for supply chain management (SCM) with a unique integration of Design for Six Sigma (DFSS) and fuzzy logic. The model helps expose important metrics and process variables, optimizes the internal works process, and recalibrates like-minded supply chain partners through joint meetings, monitoring, and proactive adjustments. The model is sufficient to ultimately address demand, surveillance, and regulation with respect to performance measurement systems. The variability of the OFP model addresses the demand for adaptability, and thus in turn, also specializes, targeting certain outcomes using appropriate metrics and process variables to incrementally optimize the overall performance.

In 2018, Kahraman et al. [14] used Data Envelopment Analysis (DEA) to assess store performance in a local Turkish retailer chain. DEA is a nonparametric programming technique that deals with multiple inputs and outputs. Inputs included store size, employee count, deliveries, and total cost, and outputs included number of customers, sales, and customer evaluations, gathered through interviews. To accommodate a changing retail environment the input-oriented BCC model with variable returns to scale was used. DEA efficiency results were analyzed using the Frontier Analyst software and provided managerial recommendations.

In 2024, Giuliana et al. [15] reported on a new inventory management model for SMEs in retail, using three tools: a dashboard-based control system for goods flow; a KANBAN canvas for process monitoring; and the PDCA (plan-do-check-act) cycle for continuous improvement. Following an initial diagnosis, inventory policies were redesigned, processes were reengineered, and employees were trained. Weekly inventory crossings between physical versus logical levels were made to track accuracy and availability.

2.2. Data Collection and Log Management

In 2017, Bradlow [16] examined the implications of big data in retail across five distinct dimensions: customers, products, time, location, and channels. It argues that enhanced data quality is driven by new sources of data, sophisticated statistical and analytical tools, widespread expertise from the field, and theoretical specification. Bayesian analysis techniques such as updating, adding to data, and hierarchical modeling are discussed in conjunction with predictive analytics and a field experimentation framework in retail applications.

In 2019, Santoro et al. [17] examined how using big data alters organizational practices and leads to benefits for retail companies. To this end, the authors were able to conduct semi-structured interviews with marketing managers from four retail companies in Italy, and they were able to look at secondary data sources to provide a broader view on how big data is used in retail operations. The study findings include practical implications and contextual understanding of the impacts of big data use, which contribute to the knowledge of big data adoption, organizational change, and improvement in firms' performance in retail businesses.

In 2018, Sirangi [18] presented a technical framework for fraud detection, which integrates log analysis in a distributed stream processing approach. The overall hybrid architecture proposes using anomaly detection, e.g., clustering and graph-based models, using an anomaly detection architecture using stream processing platforms like Apache Flink and Kafka.

In 2023, Ramos and colleagues [19] illustrated the importance of scalable monitoring, detailing key components, tools, and techniques. They stressed the use of Spring Boot Actuator for monitoring Spring Boot applications, challenges that are still evident, and areas for future expansion, including expansion into AI-driven monitoring and changes in observability.

In 2019, Castelo-Branco et al. [20] sought to develop a knowledge base for applying data mining tools and techniques in retail more effectively, including market basket analysis, association rules, and up -selling or cross-selling. Businesses integrate statistics and modeling to improve a variety of functions. The paper expands on typical business intelligence applications and provides applications of successful data mining within retail while discussing future opportunities for the data-driven approach within retail. In Table 1, several methods are demonstrated, allowing for efficient data collection and log analysis to improve retail performance.

Table 1: Methods for Data Collection and Log Analysis in Retail

Author(s)	Year	Method/Approach	Key Contribution
Bradlow [16]	2017	Big data analytics; Bayesian analysis; predictive analytics	Explored big data dimensions in retail and enhanced data quality through statistical tools and theory
Santoro et al. [17]	2019	Semi-structured interviews; secondary data analysis	Studied big data adoption and organizational transformation in Italian retail companies
Sirangi [18]	2018	Log analysis with distributed stream processing; anomaly detection (clustering, graph-based); Apache Flink & Kafka	Enhanced fraud detection in retail systems using hybrid architecture
Ramos et al. [19]	2023	Scalable monitoring framework; Spring Boot Actuator	Highlighted monitoring strategies, tools, and AI-driven observability advancements
Castelo-Branco et al. [20]	2019	Retail data mining; market basket analysis; association rules	Built knowledge base for data-driven strategies, cross-selling/up-selling, and BI applications

2.3. Metrics and Key Performance Indicators (KPI)

In 2025, Subagyo et al. [21] combined Recency, Frequency, and Monetary (RFM) analysis with K-Means clustering on transactional data from 2022–2024. RFM values were normalized using z-scores, and the optimal cluster number was determined via the Elbow Method, validated with Silhouette Score (0.52) and BSS/TSS ratio (0.80). The resulting five segments Platinum, Gold, Silver, Bronze, and B show distinct behaviors, enabling targeted marketing, loyalty programs, and reactivation campaigns.

Stoyanov (2021) [22] provided an objective measure of how well an economic operator qualifies as an agent for successful business. In more traditional retail, important performance measures include gross profit margin, operating margin, inventory turnover, return on inventory investment, and profit per employee. The paper examines the economic performance indicators of the leading side of three retail chain accounts in the fast-moving consumer goods sector, based in Bulgaria.

Radonić et al. (2016) [23] described turnover and gross margin, fundamental measures that directly contribute to profitability and provide the foundation for other measures like market share and EBIT. Much knowledge can be on the detailed analysis of these general KPI metrics, especially for lower management. Pricing strategies, such as Elvis or Every Day Low Prices (EDLP) and High-Low (Hi-Lo), can have major influences on turnover and margins. Importantly in our measure of turnover is that turnover is a function of quantity as well as price and thus good pricing strategies create even better measurement of KPIs and come in more informed decision making with aligned KPI metrics.

In 2019, Adivar et al. [24] introduced a performance management framework and roadmap for success that concentrated on comparative evaluation between traditional and omnichannel supply chains. The comprehensive indicators are grouped into four dimensions: sustainability, efficiency/effectiveness, responsiveness, and flexibility. These indicators are utilized from seven perspectives: customers, operations, sourcing, finances, IT, transportation, and environment. Two frameworks focused on customers and one statistical model indicate that physical stores are crucial for omnichannel companies to succeed.

In 2022, Luo et al. [25] presented an xDeepFM-LSTM hybrid forecasting model that enhances apparel retail companies' sales forecasting efficacy. An xDeepFM model was first applied to evaluate the sales data features' correlations, followed by forecasting sales performance. The residual errors were rectified using a Long Short-term Memory (LSTM) model to boost prediction performance.

2.4. Real-time Monitoring and Alerts

In 2024, Alghaslan [26] considered challenges related to a changeable appearance of the product, chance occlusions, lighting reflections, and retrieval with an inconsistent resource base; the process began with a YOLOv8s model to detect grocery products on two separate aisles. The second stage of recognition task was using EfficientNetV2-S and ResNet18 models, which achieved accuracy for recognition of grocery products on the shelf. Retail Eye focused on the model's use of supervised contrastive learning and adherence/matching in combining supervised learning with compliant matching which provided significant advantages in accuracy for performance compared to a one-stage thick prior approach.

Prabu [27] in 2021 created efficiencies by generating immediate insights into operations, consumer behavior, and inventories. Retailers can react quickly to changes in demand to not run out of product and avoid overstock. Real-time analytics leads to reduced costs and increases customer satisfaction through enhanced availability and quicker delivery by improving forecasting accuracy, inventory visibility, and logistics.

Sivalakshmi et al. [28] in 2024 created a smart retail security and surveillance framework that harnesses transfer learning (TL) combined with video analytics on the cloud. Cloud computing provides the ability to store video data on a large-scale and allows real-time analytics of that same data. The TL process customizes pre-trained models to execute anomaly detection accurately. The proposed framework will quickly identify suspicious activities, alert staff, and provide an avenue for dynamic scalability. While the framework is designed as a retail security platform, it also provides vendor value by conducting real-time visual analytics, including crowd analysis, heat mapping, and tracking consumer behavior. Vendors can leverage the insights to improve store layouts, product placements, and marketing strategies based on evidence-driven decisions.

Pietrini et al. [29] in 2024 proposed a Shelf Management system with deep learning integration. The system includes an object detection module (with mAP 0.752) based on RetinaNet, a deep Hough Transform model (with F1 0.97) for shelf row detection, and product recognition (with top-1 accuracy 0.93) using a MobileNetV3 model is advanced by triplet loss and a FAISS module. Two annotated datasets, SHARD and SHAPE, are introduced for accurate localization and serve as points of reference for researchers to further study optimal shelf product placement and retail automation.

Hyb-SMPC, a hybrid deep learning and computer vision-based method, was created in 2022 by Saqlain et al. [30].

Using one-stage detectors, the first module finds fine-grained retail items. The best-performing ones are chosen by comparing YOLO V4, YOLO V5, and YOLOR. By converting layouts to JSON and comparing them with processed images, the second module creates compliance reports that confirm planogram conformance.

2.5. Data Visualization and Dashboards

Surwade et al. [31] (2024) reported that, to make sense of complex information, visualizing information and data includes charts, graphs, and maps to guide working with enterprise intelligence. They enable users to quickly make sense of large sets of data, identify trends, and ultimately make decisions or plans of action. Clear and concise visuals improve communication, collaboration, and opportunities for improvement and growth which ultimately provide a competitive advantage for an organization. Visualizing data and information effectively requires consideration of design principles, user experience principles, and the organization's goals.

Rieg [32] (2025) expanded the Technology Acceptance Model (TAM) and Unified Theory of Acceptance and Use of Technology (UTAUT) with information quality (IQ) and decision quality (DQ) by IQ of lower search and reconciliation costs reduced effort and DQ improved usefulness by contributing defensible analyses. This framework brings additional information to the two frameworks as the inherent complementarity of IQ and DQ improves ease of use. The use of a dashboard improves performance based on enabling-control mechanisms. The framework is corroborated with data from German accountants and provides the ability to consider theoretical implications and practical guidance for practices in your own organization when integrating dashboards.

Sedrakyan et al. [33] (2029) located gaps between dashboards related to design and learning science concepts and proposed a conceptual model to connect design principles to learning processes and feedback types. The focus of this paper builds off that research by focusing on mapping dashboard principles to data and information visualization principles. Drawing on conceptual analysis and empirical data, this paper provides recommendations intended to guide in selection of visual representations aligned to specific types of feedback.

Although Dowding and Merrill [34] addressed general usability in 2018, they do not specifically assess systems that produce information visualizations. This research develops a domain-specific heuristic evaluation checklist in a way that combines certain of Nielsen's principles with those of earlier research around visualizations. The twelve health informatics specialists evaluated each factor using the nominal group technique via email from 1-10.

In 2017 Franklin and colleagues [35] examined the challenges and possibilities to make the emergency department (ED) real-time visualizations to support clinical decision-making. While organizational change like adding a flow coordination nurse or have a physician in triage would help overall throughput; however, these organizational policy changes do not address the immediate workflow decisions of any one clinician. Our prototype Throughput Dashboard provides a real-time view of patient status and department flow to support clinicians to make well-informed moment-to-moment decisions. Therefore, this way of thinking underlines the ability to rapidly intervene, and helps effectiveness and patient care to improve in the ED.

2.6. Integration with Retail Systems

In 2021, a conceptual framework named CECoR was introduced by Anica-Popa et al. [36] underscoring Customer Experience (CE), Cost Reduction (Co), and Revenue generation (R). Several AI applications including chatbots, smart shelves, and personalized recommendations serve to improve customer knowledge and improve operational efficiencies. The CECoR provides practical recommendations for both practitioners and researchers with clear directions on how to implement AI, manage risk, and improve benefits within retail information systems while simultaneously keeping a customer-centric focus.

In 2024, Ahmad et al. [37] explored the role of supply chain dynamism in enhancing decision-making in Jordan’s retail sector, emphasizing the moderating effects of technological and organizational integration. While supply chain dynamism is recognized as crucial for navigating today’s complex retail environment, its direct impact on decision-making has received limited attention. The research identifies key determinants that enable effective supply chain decisions and examines how collaboration, supported by technology and organizational integration, mediates this relationship.

In 2023, Chen et al. [38] examined how channel integration quality affects individual consumers' perceptions and responses. Using 517 online questionnaires collected in 2020 from Taipei (Taiwan) and Shenzhen (China) via quota sampling, structural equation modeling revealed that channel-service configuration positively influences experience quality, which in turn enhances customer engagement and empowerment. Experience quality and relationship proneness serve as mediators between channel integration quality and consumer responses. The study emphasizes a consumer-focused approach to omni-channel integration, providing insights into the mechanisms linking channel quality to customer outcomes.

In 2019, Kalisetty and Ganti [39] analyzed the impact of advancements in technology on retail and found both opportunities and challenges. The adoption of new systems can incur significant costs and includes system upgrades and protocols across sectors. Public adoption can take time, and the advantages appear to be more favorable to large retailers. Despite the challenges, big data and analytics can help predict outcomes, optimize implementation, and offer efficiencies.

In 2016, Demoulin and Djelassi [40] suggested and tested a holistic model that includes individual, system, and situational factors affecting both customers' intention to use SSTs and their actual use. this study used a survey method and collected data from 143 physical SST users and 150 non-users as they exited a grocery store. Structural equation modeling was employed to investigate relationships among constructs, complemented by logistic regression analysis to identify determinants of actual use, which helps illuminate factors that influence retail settings.

2.7. Comparative Evaluation of Tools

In 2017, Holimchayachotikul and Leksakul [41] highlighted three types of drivers: individual, system and situational drivers influence SST adoption and actual use. They collected monthly store performance data that was then cleaned, verified, and optimized through particle swarm optimization. The data was then processed using a neuro-fuzzy system capable of predicting future sales and future expenses to help managers with decisions.

In 2024, Tan et al. [42] evaluated advancements in computer vision, and its applications in retail robots through the study of its benefits to retail automation. In this work, we present a new self-checkout system using the improved YOLOv10 network with the goal of automating checkouts to use labor more efficiently and reduce overall labor costs. Our model optimized using a specialized detection head structure from YOLOv8 model results in increased accuracy of product detection during the checkout process. As an additional and independent advancement, we have developed post-processing algorithms for self-checkout scenarios that not only improve reliability but overall system performance. The framework developed contributes to future retail operations by not only advancing automated product detection but forward in the operational processes of automated processes in characterized retail store contexts.

Kanjula et al. [43], in 2022, developed a cost-efficient AI-at-Edge people counting system for measuring the daily conversion rate by monitoring visitor count and transactions. The system provided informational data from which actionable insights may improve retail operation using only the least amount of hardware and technology. Hence, the system enabled quick, and easy, but informed decision-making for store management.

Sari et al. [44], in 2024, proposed a more holistic approach embedding a formal methodology with Fuzzy Cognitive Maps (FCM) in combination with Picture Fuzzy TOPSIS (PF-TOPSIS). FCM provides the key performance criteria and assists with creating models for showing descriptive and visual complexity of relationships. FCM significantly contextualizes the underlying complexity of interrelationship dynamics between criteria and sub-criteria. PF-TOPSIS then leverages the FCM criteria to frame an evaluation for retail, alleviating uncertainty and managing context ultimately mitigate ambiguity in weighing which strategic action(s) take ranked priorities.

The research conducted by Barutbaş et al. [45] in 2024 highlights the growing challenges in evaluating businesses and the increased emphasis on valuation during and after pandemics. One of the approaches to dealing with uncertainty and vagueness is the use of decomposed fuzzy sets (DFS), which expands approaches that utilize fuzzy set logic. The application of DFS and the Decomposed Fuzzy TOPSIS (DF-TOPSIS) method, is applied to a case study of publicly listed retail firms on Borsa Istanbul (BIST). The introduced methodology incorporated important underlying concepts of identifying criteria, developing performance scoring, and aggregating performance assessments through meaningful and structured approaches under contextual uncertainty (in layman's terms uncertainty about value prospects)? The overarching premise of the methodology centers around capturing the concern of stakeholders and/or investors with limited contextual actionability for deciding, while also distinguishing between equally favorable ventures to facilitate decisions.

3. Challenges and Limitations

Retail applications also present technical, operational, and strategic issues which complicate performance optimization. The first issue is the sheer amount and speed of data produced from the retail application itself. Today's retail applications produce a staggering amount of customer engagement, transaction, and inventory data that must be collected, processed, and analyzed, almost in real time, to optimize performance. The amount of data indicates that there are considerable resource requirements for data storage, data processing, and data monitoring. Also, many retail applications are complex and composed of heterogeneous systems, including integrating point-of-sale (POS) systems and e-commerce applications, enterprise resource planning (ERP), and customer resource management (CRM) software. This heterogeneity makes end-to-end performance monitoring and optimization difficult and requires implementing secure APIs and complex integration.

Another challenge involves the ever-evolving nature of customer behavior. Customers' abrupt movements away from previous purchasing behavior, holiday shopping peaks, or limited-time promotion encounters can generate unpredictable workloads and traffic spikes that may cause latency, service outages, or a declining experience if the application stack can't maintain scalability and resilience. Similar to the previous discussion, network and infrastructure bottlenecks - such as limited bandwidth, server failures, or cloud service outages - will impact an application's performance irrespective of any optimizing at the software level. Many monitoring tools, like Splunk, Datadog, and New Relic, provide deep visibility into application health, but each has its limitations. For example, Splunk can become prohibitively expensive when dealing with large scale log data, Datadog is more expensive when using extensive multi-service monitoring, and New Relic has limited log management capabilities which may require other tools to compensate. Further, alert fatigue is its own limitation to operations; excessive alerts that are poorly configured can create an overwhelming situation for IT teams, resulting in delayed response times, or critical incidents being missed altogether. In addition, security and privacy limitations exist when considering optimization strategies, as retail applications often require the handling of sensitive financial and personal data according to regulations like PCI DSS and GDPR. In addition to combining high performance with stringent compliance requirements, there is ongoing cost vs performance trade-offs, as optimal performance will often require significant investments in infrastructure as well as also requiring advanced monitoring and optimization solutions, which can conflict with budgetary restrictions. Although real-time monitoring has advanced, the predictive power of existing technologies is still restricted, making it harder to proactively stop performance decline before consumers are impacted. When taken as a whole, these difficulties show how difficult it is to maintain the best possible performance for retail applications and emphasize the necessity of comprehensive, flexible, and astute performance management techniques.

4. Case Studies

It is of utmost importance to optimize the performance of retail applications to improve customer experience and increase sales. Modern monitoring solutions such as Splunk, Datadog, and New Relic allow retailers to see system metrics, identify bottlenecks, and solve problems proactively. Below are case studies that demonstrate how these three tools have been employed in cases related to an e-commerce platform, a POS system, a mobile app, and in omnichannel retail that produced relevant performance results.

E-Commerce Platform Latency Reduction

A large online retailer had a critical problem with a slow checkout process that was resulting in high cart abandonment. The retailer implemented Datadog to monitor live real-time transactions across the platform. Alerts were enabled to notify if the calls to the payment gateway were slow. Using root cause analysis to monitor the issue, database queries were eventually identified as blocking the process. The retailer reduced their checkout latency by 40% and increased the number of completed transactions by 15%, which serviced customer satisfaction.

Point-of-Sale (POS) System Uptime Improvement

A retail chain with several brick-and-mortar locations was experiencing periodic crashing of its POS systems, most often at peak hours, that interrupt the sales process.

The company expanded the use of Splunk to collect server logs and analyze the application errors. The main reason was determined to be recurrent memory leaks using anomaly detection software. The retailer greatly decreased downtime after resolving these problems, increasing POS uptime to 99.9%. To ensure seamless shop operations and improved customer service, issue resolution times decreased from an average of three hours to only thirty minutes.

Improving Customer Experience with Application Insights

An online grocery store was experiencing slow performance issues with their mobile application, resulting in customer complaints and lower engagement with their app. The team decided to utilize the New Relic APM service to gain visibility into the mobile application and backend APIs. During monitoring, they discovered that the calls to the inventory API were causing users to experience slow performance, so the development team made some code improvements and caching adjustments to enhance the speed of the APIs. As a result of these actions, the load times for the application were reduced by 50%, which led to overall improved customer satisfaction, retention and repeat purchases.

Optimizing Retail Inventory Management

A fashion retail chain was experiencing slow inventory updates in their online store, which negatively impacted product availability and sales. They utilized Splunk dashboards to continuously monitor the backend inventory update services to identify the slow updates happening in the batch service. They set alerts for updates that were not updated within their normal thresholds so that the team could proactively address issues. The team also utilized predicted analysis to forecast high-traffic periods so product availability could be managed. This initiative resulted in reducing stock-outs by 30% and improving sales forecasts accuracy, and ultimately better managing their inventory.

Tracking Performance Across Channels

One omni-channel retailer was experiencing inconsistent performance between their website and mobile app and their in-store kiosks, resulting in a poor customer experience. They brought in Datadog to enable centralized monitoring across all channels with end-to-end tracing so the IT team could identify latency between services. With databases unified into customizable dashboards, the IT team was able to identify and resolve cross-channel performance issues more effectively. This resulted in 25% faster resolution of performance bottlenecks and a smooth, reliable customer shopping experience across all retail channels.

5. Discussion

In a competitive environment, retail application performance optimization is essential due to the huge impact on revenue related to customer experience and operational effectiveness. The analysis of tools such as Splunk, Datadog, and New Relic is an important demonstration of how contemporary monitoring tools have enabled proactive performance management across several retail contexts. Each tool has unique advantages. Splunk demonstrates superior log aggregation and search, Datadog provides full monitoring on application and infrastructure and has perfect cloud integration, and New Relic offer both insight into application performance and real-time analytics. As noted in this analysis, retail applications demonstrate value because of both metric based monitoring and real-time alerts. For example, attention to transaction latency, error rates, and server resource monitoring will help retailers identify performance issues or bottlenecks proactively versus reactively with help from data. Automated alerts and anomaly detection have the potential to bolster application down time and service redundancy on high-volume traffic events such as retail sales or during the holiday seasons for instance.

Performance optimization benefits significantly from data visualization and dashboards. Making relevant KPIs intuitive and actionable increases the speed and quality of retail managers and developers' decisions. Splunk's engine for data analytics, Datadog's dashboard focused on infrastructure, and New Relic's dashboards focused on applications all provide actionable insights for performance, operational efficiency, and customer satisfaction. However, there can be complications when it comes to the selection and implementation of the monitoring tool. The analysis could be impeded by complexity of integration, cost, scalability, and readiness of the organization for a performance optimization project. In addition, while these tools can provide useful insights, the analysis process relies on a human element - to identify, analyze, design optimization strategies, and ensure alignment with the business. Future research and practice need to focus on hybrid monitoring strategies that combine real-time analytics with predictive modeling and AI/ML strategies to enhance sensitivity to issues before they arise without resource drain.

Incorporating these strategies could help with ongoing improvements to performance, resource allocation, and pre-emptively detecting issues for retail applications. In closing, strategically using performance monitoring tools and support for data-driven decisions can greatly enhance retail systems efficiency, reliability, and responsiveness.

The results of this review point to several exciting potential areas for future research into optimizing application performance for retail (and elsewhere). While current state-of-the-art monitoring tools provide real-time observability, the increasingly complex nature of the retail ecosystem and all it entails indicates a need for predictive, intelligent, and adaptive mechanisms to observe, detect, and manage Application Performance Management (APM) in the years ahead.

AI-Powered Predictive Analytics

Future research should examine the use of machine learning models, such as recurrent neural networks (RNN's), transformers, and ensemble models to predict workload spikes, transaction anomalies, and bottlenecks in the infrastructure. Predictive models enable mitigation and avoidance of downtime and enhance the consumer experience through scaling and the avoidance or proactive resolution of faults.

Automated Root Cause Analysis (RCA)

Most current monitoring solutions can create alerts but do not provide sufficient diagnostic information. Use of explainable Artificial Intelligence (XAI) within current observability frameworks can provide automated RCA as the system correlates system metrics, consumer behavior, and log traces, resulting in increased operational efficiency and reduction of mean-time-to-resolution (MTTR).

Adaptive Optimization via Reinforcement Learning

Retail workloads typically experience sudden changes in a work pattern during promotional events, the holiday season, and high traffic periods. In multi-cloud or hybrid environments, deployment of reinforcement learning agents to dynamically optimize cloud resource allocation, to provide a balance between performance reliability and saving on costs, would be a productive route of future research.

Cross-Channel Predictive Monitoring

With the proliferation of omnichannel retail, monitoring systems must be designed to encompass e-commerce platforms, mobile apps, physical stores, and Internet of Things (IoT) retail devices. Future research could design unified predictive monitoring architecture that allows for seamless service delivery across these disparate environments.

Security-Aware Performance Monitoring

As compliance regulations such as Payment Card Industry Data Security Standard (PCI DSS) and General Data Protection Regulation (GDPR) become more stringent, future frameworks for monitoring should include security anomaly detection as well as performance metrics. Joint optimization models can facilitate performance and security that optimizes throughput, latency, and compliance without compromise.

Hybrid Cloud and Edge Intelligence

Retail systems are increasingly integrating centralized cloud infrastructure with decentralized edge devices (e.g. smart shelves, kiosks, and point-of-sale terminals). Adding predictive monitoring at the edge enables immediate anomaly detection and localized response, which can provide feedback into centralized dashboards to provide a global view.

6. Conclusion

The authors of this paper conducted a systematic review and analysis of optimization strategies, monitoring frameworks, and best practices aimed at retail application performance. The review identified 45 peer-reviewed, research articles published in high-quality journals and conferences in the time frame from 2015 to 2025 with an aim of establishing methods, limitations, and areas of research. The review covered areas including performance monitoring techniques and tools, data collection and log management, key performance indicators (KPIs), real-time monitoring, data visualizations, and integration into existing systems. In addition, this study argued about the relative effectiveness of monitoring platforms, the limitations of current optimization practices, and the promise of intelligent analytics for enhancing resilience and scalability. The resulting findings provided academic and practice-based insights for researchers and practitioners who seek to design, deploy, and optimize the performance of retail applications at scale in cloud-based computing environments.

References

1. Aghakarimi, E., Karimi, H., Aghsami, A. and Jolai, F., 2024. Evaluating and improving the performance of retailers' branches by considering resilience, sustainability and sales-marketing. International Journal of Productivity and Performance Management, 73(5), pp.1357-1386.
2. Barasin, A.M., Alqahtani, A.Y. and Makki, A.A., 2024. Performance evaluation of retail warehouses: A combined MCDM approach using G-BWM and RATMI. Logistics, 8(1), p.10.
3. Santamaria-Ruiz, M., Ortíz-Morales, M., Benavides-Morón, E., Vargas-Suarez, E. and Troncoso-Palacio, A., 2023. An analysis of management by competencies in grocery retail distributors. Procedia Computer Science, 224, pp.431-436.
4. Kawa, A. and Zdrenka, W., 2024. Logistics value in e-commerce and its impact on customer satisfaction, loyalty and online retailers' performance. The International Journal of Logistics Management, 35(2), pp.577-600.
5. Nazir, H. and Fan, J., 2024. Revolutionizing retail: Examining the influence of blockchain-enabled IoT capabilities on sustainable firm performance. Sustainability, 16(9), p.3534.
6. Lagzi, M.D. and Taghizadeh-Yazdi, M., 2024. A hybrid stochastic data envelopment analysis and decision tree for performance prediction in retail industry. Journal of Retailing and Consumer Services, 80, p.103908.
7. Prashar, A., 2023. Supply chain sustainability drivers for fast-moving consumer goods (FMCG) sector: an Indian perspective. International Journal of Productivity and Performance Management, 72(8), pp.2397-2419.
8. Zkik, K., Belhadi, A., Kamble, S., Venkatesh, M., Oudani, M. and Sebbar, A., 2024. Cyber resilience framework for online retail using explainable deep learning approaches and blockchain-based consensus protocol. Decision Support Systems, 182, p.114253.
9. Banerjee, S. and Parisa, S.K., 2023. AI-Enhanced Intrusion Detection Systems for Retail Cloud Networks: A Comparative Analysis. Transactions on Recent Developments in Artificial Intelligence and Machine Learning, 15, p.15.
10. Lu, W.M., Kuo, K.C. and Tran, T.H., 2023. Impacts of positive and negative corporate social responsibility on multinational enterprises in the global retail industry: DEA game cross-efficiency approach. Journal of the Operational Research Society, 74(4), pp.1063-1078.
11. Harrauer, V. and Schnedlitz, P., 2016. Impact of environment on performance measurement design and processing in retailing. International Journal of Retail & Distribution Management, 44(3).
12. Parisa, S.K., Banerjee, S. and Whig, P., 2023. AI-Driven Zero Trust Security Models for Retail Cloud Infrastructure: A Next-Generation Approach. International Journal of Sustainable Devlopment in field of IT, 15, p.15.
13. Chen, H., Heydari, M., Lai, K.K. and Zhang, J., 2025. Enhancing supply chain resilience in retail operations: a novel DFSS and fuzzy logic model for optimizing order fulfillment process. Annals of Operations Research, pp.1-31.
14. Kahraman, C., Uluğ, İ., Othan, C.B., Özkan-Özen, Y.D. and Kazançoğlu, Y., 2018, August. Efficiency analysis in retail sector: Implementation of data envelopment analysis in a local supermarket chain. In The International Symposium for Production Research (pp. 884-897). Cham: Springer International Publishing.
15. Giuliana, N.P., Panizo, P., Zadott, T., Alvarez, O., Nicolas, R. and Ortiz, M., 2024, October. Increase inventory accuracy and availability through the implementation of a logistics process control model based on lean tools in SMEs in the retail sector. In 2024 9th International Engineering, Sciences and Technology Conference (IESTEC) (pp. 1-6). IEEE.
16. Bradlow, E.T., Gangwar, M., Kopalle, P. and Voleti, S., 2017. The role of big data and predictive analytics in retailing. Journal of retailing, 93(1), pp.79-95.
17. Santoro, G., Fiano, F., Bertoldi, B. and Ciampi, F., 2019. Big data for business management in the retail industry. Management Decision, 57(8), pp.1980-1992.
18. Sirangi, A., 2018. Retail Fraud Detection via Log Analysis and Stream Processing. Computer Fraud & Security Bulletin, 2018, pp.21-32.
19. Ramos, A., Scalable Monitoring Solutions for Enterprise Applications.
20. [20] Castelo-Branco, F., Reis, J.L., Vieira, J.C. and Cayolla, R., 2019. Business intelligence and data mining to support sales in retail. In Marketing and Smart Technologies: Proceedings of ICMarkTech 2019 (pp. 406-419). Singapore: Springer Singapore.

21. Subagyo, U., Thoha, A.B., Syafrianto, A., Santoso, H. and Sanuri, R., 2025, June. Customer Behavior Profiling in Wholesale Retail Using RFM Analysis and K-Means Clustering. In 2025 4th International Conference on Electronics Representation and Algorithm (ICERA) (pp. 467-472). IEEE.

22. Stoyanov, M., 2021. Measuring Key Performance Indicators in Retail Trade. Izvestiya. Journal of Varna University of Economics, 65(2), pp.161-175.

23. Radonić, M., Đorđević, V. and Omerbegović-Bijelović, J., Key Performance Indicators and Pricing Strategies Planning in Retail Industry. OPERATIONS MANAGEMENT 946, p.995.

24. Adivar, B., Hüseyinoğlu, I.Ö.Y. and Christopher, M., 2019. A quantitative performance management framework for assessing omnichannel retail supply chains. Journal of Retailing and Consumer Services, 48, pp.257-269.

25. Luo, T., Chang, D. and Xu, Z., 2022. Research on Apparel Retail Sales Forecasting Based on xDeepFM-LSTM Combined Forecasting Model. Information, 13(10), p.497.

26. Alghaslan, M., Almutairy, K., El-Alfy, E.S. and Siddiqui, A.J., 2024, December. RetailEye: Supervised Contrastive Learning with Compliance Matching for Retail Shelf Monitoring. In International Conference on Neural Information Processing (pp. 248-263). Singapore: Springer Nature Singapore.

27. Prabu, V.P., The Impact of Real-Time Data Analytics on Retail Supply Chain Efficiency.

28. Sivalakshmi, P., Kavitha, U., Usha, R., Pattanaik, O., Maniraj, S.P. and Srinivasan, C., 2024, August. Smart retail store surveillance and security with cloud-powered video analytics and transfer learning algorithms. In 2024 second International conference on intelligent cyber physical systems and internet of things (ICoICI) (pp. 242-247). IEEE.

29. Pietrini, R., Paolanti, M., Mancini, A., Frontoni, E. and Zingaretti, P., 2024. Shelf Management: A deep learning-based system for shelf visual monitoring. Expert Systems with Applications, 255, p.124635.

30. Saqlain, M., Rubab, S., Khan, M.M., Ali, N. and Ali, S., 2022. Hybrid Approach for Shelf Monitoring and Planogram Compliance (Hyb-SMPC) in Retails Using Deep Learning and Computer Vision. Mathematical Problems in Engineering, 2022(1), p.4916818.

31. Surwade, N.B., Shiragapur, B. and Hussain, A., 2024. Data visualization and dashboard design for enterprise intelligence. In Metaheuristics for Enterprise Data Intelligence (pp. 71-91). CRC Press.

32. Rieg, R., 2025. Exploring the determinants and performance effects of digital dashboard use by management accountants. Journal of Management Control, pp.1-36.

33. Sedrakyan, G., Mannens, E. and Verbert, K., 2019. Guiding the choice of learning dashboard visualizations: Linking dashboard design and data visualization concepts. Journal of Computer Languages, 50, pp.19-38.

34. Dowding, D. and Merrill, J.A., 2018. The development of heuristics for evaluation of dashboard visualizations. Applied clinical informatics, 9(03), pp.511-518.

35. Franklin, A., Gantela, S., Shifarraw, S., Johnson, T.R., Robinson, D.J., King, B.R., Mehta, A.M., Maddow, C.L., Hoot, N.R., Nguyen, V. and Rubio, A., 2017. Dashboard visualizations: Supporting real-time throughput decision-making. Journal of biomedical informatics, 71, pp.211-221.

36. Anica-Popa, I., Anica-Popa, L., Rădulescu, C. and Vrîncianu, M., 2021. The integration of artificial intelligence in retail: benefits, challenges and a dedicated conceptual framework. Amfiteatru Economic, 23(56), pp.120-136.

37. Ahmad, A.Y.B., Allahham, M., Almajali, W.I., Ayasrah, F.T. and Sabra, S., 2024, December. Supply Chain Innovation on Acceleration Decision-Making, The Mediating Role of Tech and Integration in the Retail Sector. In 2024 25th International Arab Conference on Information Technology (ACIT) (pp. 1-6). IEEE.

38. Chen, T.Y., Yeh, T.L., Wu, H.L. and Deng, S., 2023. Effect of channel integration quality on consumer responses within omni-channel retailing. Asia Pacific Journal of Marketing and Logistics, 35(1), pp.149-173.

39. Kalisetty, S. and Ganti, V.K.A.T., 2019. Transforming the Retail Landscape: Srinivas's Vision for Integrating Advanced Technologies in Supply Chain Efficiency and Customer Experience. Online Journal of Materials Science, 1, p.1254.

40. Demoulin, N.T. and Djelassi, S., 2016. An integrated model of self-service technology (SST) usage in a retail context. International Journal of Retail & Distribution Management, 44(5), pp.540-559.

41. Holimchayachotikul, P. and Leksakul, K., 2017. Predictive performance measurement system for retail industry using neuro-fuzzy system based on swarm intelligence. Soft Computing, 21(7), pp.1895-1912.

42. Tan, L., Liu, S., Gao, J., Liu, X., Chu, L. and Jiang, H., 2024. Enhanced self-checkout system for retail based on improved YOLOv10. Journal of Imaging, 10(10), p.248.

43. Kanjula, K.R., Reddy, V.V. and Abraham, J.S., 2022. People counting system for retail analytics using edge AI. arXiv preprint arXiv:2205.13020.

44. Sari, I.U., Ozkirim, C. and İlgin, U., 2024, July. Integrating Fuzzy Cognitive Mapping and Picture Fuzzy TOPSIS for Long-Term Retail Performance Evaluation. In International Conference on Intelligent and Fuzzy Systems (pp. 630-637). Cham: Springer Nature Switzerland.

45. [Barutbaş, K., Bulut, B. and Otay, I., 2024, July. Financial performance evaluation of leading retail companies in BIST utilizing decomposed fuzzy TOPSIS. In International Conference on Intelligent and Fuzzy Systems (pp. 123-132). Cham: Springer Nature Switzerland.