Sarcouncil Journal of Applied Sciences

ISSN(Online): 2945-3437

Volume- 06| Issue- 04| 2026

Research Article

Received: 15-03-2026 | **Accepted:** 10-04-2026 | **Published:** 25-04-2026

Real-Time Supply Chain Visibility Using Kafka-Based Data Streaming Architectures

Rishabh Agarwal

Harrisburg University of Science and Technology, Pennsylvania.

Abstract: Real-time supply chain visibility has emerged as a critical enabler for operational excellence, risk mitigation, and rapid decision-making in modern enterprises. With global supply chains becoming increasingly dynamic, interconnected, and data-intensive, traditional batch-processing systems are no longer sufficient to handle the velocity and variability of supply chain events. Kafka-based data streaming architectures offer a scalable, fault-tolerant, and high-throughput foundation for continuously processing logistics movements, inventory fluctuations, transportation updates, and supplier interactions as they occur. This article examines the fundamental principles, architectural considerations, and practical applications of Kafka-driven real-time supply chain visibility. It further discusses advanced analytics, predictive modeling, and integration patterns powered by streaming data flows. Emphasis is placed on Kafka's distributed log structure, pub-sub mechanisms, event streaming pipelines, and stream processing capabilities, and how these collectively address long-standing challenges such as data latency, fragmented visibility, silos across suppliers, inconsistent system updates, and limited orchestration across global supply chain networks. The article concludes by highlighting the transformative impact of Kafka-enabled architectures in enhancing resilience, transparency, and decision intelligence across end-to-end supply chain ecosystems.

Keywords: Supply Chain , Architectur, Risk mitigation.

INTRODUCTION

Global supply chains are evolving into highly complex and interconnected systems that generate massive volumes of data from logistics networks, transportation fleets, manufacturing lines, warehouses, suppliers, and retail endpoints. Traditional supply chain management approaches often rely on static, batch-oriented systems that cannot provide the level of immediacy, granularity, and continuity required for predictive, data-driven decision-making. Real-time visibility has therefore become indispensable, enabling organizations to identify disruptions early, optimize inventory movements, respond to demand variations, and ensure compliance across geographically dispersed operations [Somapa, S. *et al*., 2018; Machado, D. S. M. *et al*., 2021; Koberg, E., & Longoni, A. 2019]. The emergence of distributed streaming platforms such as Apache Kafka presents a powerful opportunity to process and integrate continuous data flows at scale, offering a robust infrastructure for capturing events at sub-second latency and transforming them into actionable intelligence [Isah, H. *et al*., 2019; Bousdekis, A. *et al*., 2021; Michelson, B. M. 2006]. This paradigm shift toward event-driven supply chain architectures is further accelerated by digital transformation initiatives, IoT proliferation, and the adoption of cloud-native technologies that emphasize speed, agility, and elasticity [Tavana, M. *et al*., 2022; Ameer, T., & Valilai, O. F. 2025; Saleh, H. H. *et al*., 2024]. Real-time supply chain visibility entails the continuous gathering, harmonization, and interpretation of data across multiple nodes, manufacturing plants, distribution centers, carriers, and retail interfaces, thereby eliminating blind spots caused by legacy connectivity and slow data propagation [Islam, A. M. M. S. 2025; Tian, X. *et al*., 2015; Sim, C. *et al*., 202]. Kafka's distributed event log and pub–sub model provide an ideal backbone for ensuring that these data streams remain reliable, scalable, fault-tolerant, and capable of serving multiple consumers simultaneously [Balakrishnan, M. *et al*., 2013; Carzaniga, A. *et al*., 2000; Cristian, F. 1991]. As organizations adopt advanced analytics, machine learning, and digital twins to enhance forecasting precision, the role of high-quality, time-sensitive data becomes even more pronounced. Kafka-based architectures ensure that analytics engines and operational systems receive uninterrupted feeds of well-structured data, empowering them to generate actionable insights for decision-makers. Since the introduction establishes the challenge of real-time visibility and the need for streaming platforms, it naturally leads to a deeper exploration of how the foundational characteristics of Kafka support these goals. Therefore, the next section transitions into a comprehensive overview of Kafka's architectural features and their relevance to real-time supply chains.

***Corresponding Author:** Rishabh Agarwal

DOI- https://doi.org/10.5281/zenodo.20737269

Kafka as the Foundation for Real-Time Supply Chain Visibility

Building on the need for seamless and continuous data integration discussed in the introduction, Apache Kafka emerges as an architectural cornerstone capable of supporting the velocity and variability of supply chain data as shown in Figure 1. Kafka's distributed log structure enables event ordering, durability, and replayability, which are essential properties for applications that depend on accurate and time-sensitive logistics information. Supply chains frequently encounter fragmented data due to the heterogeneity of partners, systems, and devices; Kafka addresses this through its robust pub-sub mechanism, where producers generate events such as shipment locations or inventory adjustments, and multiple consumers independently process them for tracking dashboards, anomaly detection, planning engines, or compliance applications [Javed, M. H. *et al*., 2017; Bruckner, R. M. *et al*., 2002; Majumder, A. *et al*., 2009]. Kafka's inherent scalability allows supply chains to support millions of events per second without compromising system performance. This is particularly important for IoT-enabled environments where sensors on vehicles, storage facilities, and handling equipment continuously transmit telemetry readings. Kafka's partitioning mechanism supports horizontal scalability, enabling high levels of parallelism for data ingestion and stream processing. Additionally, Kafka's ability to handle heterogeneous data formats, from GPS coordinates to ERP updates to RFID signals, enhances the integrability of supply chain systems, ensuring that events from different geographic or organizational domains converge into a unified and accessible real-time data pipeline [Elsaleh, T. *et al*., 2020; Wenzel, H. *et al*., 2019; Lekić, M. *et al*., 2021].

Another critical capability is Kafka's durability and fault tolerance. As global supply chains operate across time zones and geopolitical boundaries, resilience becomes non-negotiable. Kafka replicates data across brokers to ensure high availability, minimizing the risk of disruptions caused by hardware failures or network outages. Furthermore, Kafka's replay capability allows downstream consumers to reconstruct event histories for root-cause analysis or regulatory audits, enhancing the transparency and traceability of supply chain activities. These foundational attributes make Kafka an excellent candidate for enabling visibility that is both real-time and reliable. Having described Kafka's core strengths, it becomes essential to contextualize how these architectural capabilities directly translate into operational improvements within end-to-end supply chain networks. This necessitates a deeper examination of visibility requirements and event-driven integration across procurement, production, logistics, and distribution processes.

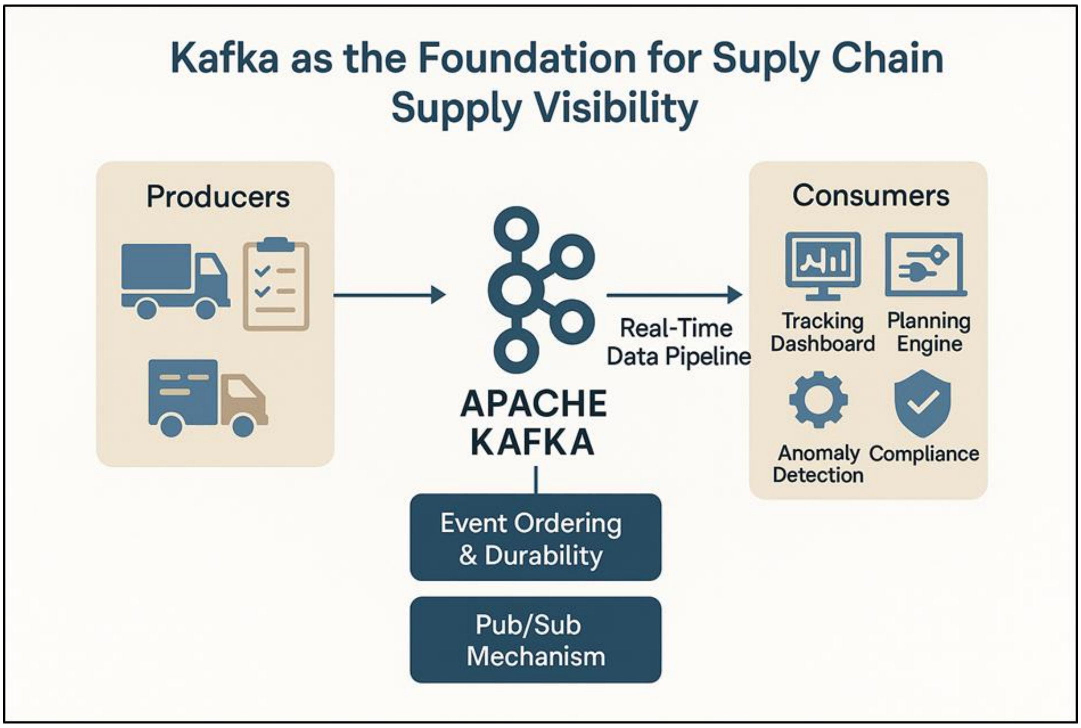

Figure 1: Apache Kafka enables real-time supply chain visibility by connecting producers and consumers through a scalable, fault-tolerant event streaming platform that supports high-throughput data pipelines.

Publisher: SARC Publisher

Real-Time Supply Chain Visibility Requirements and Kafka's Role

Following the exploration of Kafka's technical capabilities, this section shifts from architectural components to functional requirements within supply chain management. Real-time visibility requires continuous synchronization of multiple data sources such as warehouse management systems, transportation monitoring tools, supplier platforms, and manufacturing execution systems. Kafka facilitates this synchronization by enabling low-latency event flows that mirror actual operational states instead of relying on periodic updates. These continuous event flows help organizations overcome the challenge of data staleness, which historically hindered responsiveness and agility in supply chain execution [Hahn, C. K. *et al*., 2000; Van der Aa, H. *et al*., 2015; Boulos, P. F. *et al*., 1994]. Visibility encompasses several dimensions, including inventory accuracy, shipment tracking, production progress, demand fluctuations, and supplier performance. Kafka's event-driven model supports each of these dimensions by allowing systems to publish and subscribe to dedicated data streams. For example, RFID scans at docking stations generate inbound and outbound inventory events, GPS devices on trucks publish location updates, while sensors in cold chain logistics send temperature readings. These events flow into Kafka topics that various analytics, monitoring, and decision-making tools consume to produce near real-time dashboards and alerts [Cristian, F. 1991; Javed, M. H. *et al*., 2017].

Traditional supply chain systems often suffer from siloed data sources that do not communicate in real time. Kafka overcomes this fragmentation by acting as a central event distribution layer that decouples producers and consumers. This decoupling ensures that disruptions in one part of the system do not cascade across the entire supply chain ecosystem. Additionally, Kafka's support for schema evolution ensures that supply chain data formats can change over time without breaking downstream applications, which is crucial for long-term scalability in complex multi-stakeholder environments [Elsaleh, T. *et al*., 2020]. With these visibility requirements established, it becomes clear that real-time supply chain execution depends not just on data movement but also on the ability to analyze and transform event streams with intelligence. Therefore, the following section transitions into real-time processing, stream transformations, and analytics capabilities supported by Kafka-based ecosystems.

To further contextualize the role of Kafka in addressing diverse visibility dimensions, Table 1 provides a comparative overview of traditional supply chain systems versus Kafka-enabled architectures across core operational metrics.

Table 1: Comparative Analysis of Traditional Supply Chain Systems vs Kafka-Based Architectures

Visibility Dimension	**Traditional Systems**	**Kafka-Enabled Architecture**
Data Latency	High – periodic batch updates (e.g., every few hours)	Sub-second – real-time streaming of events
Inventory Tracking	Fragmented across WMS, ERP, and retail platforms	Unified through continuous event ingestion from all inventory sources
Transport Monitoring	Manual updates, GPS polling in intervals	Continuous GPS event streaming with predictive analytics
Supplier Integration	Point-to-point integration; slow EDI processes	Event-driven, decoupled communication via Kafka topics
Analytics and Forecasting	Dependent on historical data; delayed insights	Real-time models updated continuously with fresh data
Scalability	Limited by central systems and rigid architectures	Horizontally scalable through Kafka partitions and distributed consumers
Resilience and Fault Tolerance	Central points of failure	Distributed brokers and replicated logs ensure failover and high availability

This comparative framework illustrates that Kafka's architectural flexibility and real-time capabilities directly align with modern supply chain visibility requirements, laying the groundwork for intelligent automation and agile responsiveness.

Stream Processing and Advanced Analytics in Kafka-Driven Supply Chains

As the previous section emphasized the need for real-time visibility, the natural extension is understanding how analytics engines transform raw event streams into actionable insights. Kafka

Publisher: SARC Publisher

supports this through its integrated stream processing components such as Kafka Streams and ksqlDB, which enable continuous transformations, aggregations, joins, and pattern detections across event streams. These capabilities eliminate the latency associated with batch-oriented ETL systems, thereby allowing organizations to identify anomalies, predict disruptions, and optimize resource movements as events occur [Lekić, M. *et al*., 2021; Boulos, P. F. *et al*., 1994]. Real-time analytics play a pivotal role in supply chain orchestration. Predictive models built on streaming data can forecast stockouts, detect route deviations, predict lead-time variability, and anticipate equipment failures. Kafka-powered pipelines ensure that these models remain continuously updated with fresh data, improving their accuracy and relevance. The combination of real-time ingestion and real-time analytics unlocks new possibilities for proactive supply chain management, such as dynamic rerouting of shipments, automated replenishment triggers, intelligent allocation of warehouse resources, and real-time quality assurance in manufacturing.

Event-driven analytics also bring value to compliance monitoring and sustainability efforts. Stream processing can detect temperature breaches in perishable goods, carbon emission patterns in transportation fleets, or deviations in production parameters that may impact product quality or regulatory adherence. Kafka's distributed architecture ensures that even computationally intensive analytics workloads do not interfere with the ingestion layer, maintaining system stability under high throughput scenarios [Hahn, C. K. *et al*., 2000]. To fully capitalize on real-time analytics, Kafka pipelines must be integrated with a broader ecosystem including machine learning platforms, cloud-native microservices, and operational dashboards. This leads to the next section, which discusses the architectural design patterns necessary to connect Kafka pipelines with enterprise systems in a seamless and scalable manner.

Kafka-Based Architectural Patterns for Supply Chain Integration

Since the previous section illustrated the analytics potential of streaming pipelines, it is essential to explore how these pipelines integrate into broader enterprise architectures. Kafka's role extends beyond data transportation; it becomes the backbone of event-driven supply chain ecosystems through well-defined integration patterns. Common patterns include event streaming hubs, CQRS-based architectures, microservices integration, and event-sourcing models.

Kafka can serve as a central nervous system connecting ERP, WMS, TMS, MES, IoT platforms, and supplier systems. Through connectors and integration frameworks, Kafka enables bidirectional interaction between legacy systems and modern cloud-native components. Suppliers can publish ASN updates, carriers send dynamic routing information, and warehouses stream cycle count events to a unified Kafka platform, enabling complete end-to-end visibility [Tian, X. *et al*., 2015; Sim, C. *et al*., 2022]. Microservices architectures benefit significantly from Kafka, as services can communicate asynchronously, improving system resilience and reducing coupling. For instance, an order management service may listen to inventory availability events, while a logistics microservice publishes route optimization updates. This modular event-driven structure enhances scalability and reduces integration overhead.

Another important pattern is event sourcing. Kafka's immutable log makes it suitable for storing event histories that represent supply chain transactions, supporting auditability and system reconstruction. This is crucial for industries that require strict traceability, such as pharmaceuticals, food processing, and high-value electronics. As we understand these integration patterns, it becomes necessary to examine how Kafka-based visibility systems translate into tangible improvements in operations. Therefore, the next section focuses on real-world use cases demonstrating the impact of Kafka architectures across various supply chain functions [Cristian, F. 1991].

PRACTICAL USE CASES AND APPLICATIONS OF KAFKA IN SUPPLY CHAINS

Building on the integration patterns outlined previously, this section highlights how Kafka-enabled real-time visibility manifests in operational scenarios. One common use case is real-time shipment tracking, where GPS devices, telematics systems, and mobile applications continuously stream position updates into Kafka topics. This supports dynamic ETA predictions, automated exception handling, and continuous route optimization. Inventory visibility also benefits from Kafka-powered architectures. Warehouse sensors, forklifts, conveyor belts, and barcode scanners produce real-time inventory events, enabling synchronized stock levels across

Publisher: SARC Publisher

systems. This eliminates discrepancies that lead to stockouts, overstocks, and fulfilment delays. Real-time visibility also improves demand forecasting by providing accurate sell-through data from retail channels.

Manufacturing operations leverage Kafka to monitor production-line sensors, machine health indicators, and quality control events. Real-time analytics detect deviations and enable predictive maintenance, reducing downtime and scrap rates. Kafka also supports multi-tier visibility, enabling transparency across suppliers and contract manufacturers. Cold chain logistics is another critical use case. Temperature, humidity, and vibration sensors continuously transmit environmental data into Kafka, allowing supply chain managers to detect breaches and proactively mitigate spoilage risks. Event triggers can initiate automated corrective actions such as alerting drivers, adjusting refrigeration units, or rerouting shipments [Javed, M. H. *et al*., 2017; Bruckner, R. M. *et al*., 2002; Majumder, A. *et al*., 2009; Elsaleh, T. *et al*., 2020].

As these examples demonstrate the operational value of Kafka architectures, the next section explores the challenges organizations face in deploying Kafka-based visibility systems and strategies to overcome them.

To demonstrate Kafka's diverse impact across sectors, Table 2 outlines practical use cases of Kafka-enabled visibility systems categorized by industry, along with the corresponding benefits achieved.

Table 2: Kafka Use Cases in Supply Chain Visibility Across Different Industries

Industry	**Kafka Use Case**	**Benefits Achieved**
Retail	Real-time stock level updates from POS and warehouses	Minimized stockouts, improved replenishment accuracy
Manufacturing	Machine sensor streaming and predictive maintenance	Reduced downtime, improved OEE (Overall Equipment Effectiveness)
Pharmaceuticals	Cold chain temperature monitoring and alert automation	Enhanced compliance, reduced spoilage of temperature-sensitive medicines
Automotive	Component tracking and supply synchronization with vendors	Just-in-time inventory, improved supplier collaboration
Logistics	Live vehicle tracking and ETA recalculation via Kafka	Dynamic route optimization, reduced delays
Food & Beverage	Quality control via sensor-based event streams	Real-time defect detection, reduced recalls
E-commerce	Customer order tracking and live delivery notifications	Improved customer experience, higher satisfaction scores

These use cases exemplify Kafka's cross-industry relevance, with each implementation tailored to address specific visibility and responsiveness challenges, thereby reinforcing Kafka's value as a strategic enabler of operational transformation.

Challenges, Limitations, and Solutions in Kafka-Based Supply Chain Architectures

Since the prior section highlighted practical benefits, it is necessary to address operational challenges associated with deploying Kafka at scale. Real-time supply chain visibility requires handling heterogeneous data formats, integrating with legacy systems, and managing governance across multiple stakeholders. Kafka's flexibility helps address many issues, but organizations often face challenges in schema evolution, data quality management, and ensuring consistent security across distributed environments. Scalability also introduces complexities. As event volumes grow exponentially due to IoT and global expansion, organizations must adopt strategies for topic partitioning, consumer group optimization, and cluster scaling. Kafka's distributed nature mitigates bottlenecks, but improper planning can result in increased latency or uneven load distribution [Van der Aa, H. *et al*., 2015; Boulos, P. F. *et al*., 1994]. Data governance and security pose additional concerns. Supply chain data often includes sensitive information about manufacturing processes, inventory levels, and transportation routes. Ensuring encryption, access control, data masking, and compliance with regulatory frameworks is essential. Kafka offers features such as ACLs, encryption, and schema validation, but organizations must implement comprehensive governance frameworks.

Another challenge is integrating Kafka with downstream systems that may not be natively

Publisher: SARC Publisher

event-driven. Middleware layers, connectors, and stream processing frameworks help bridge these gaps but require careful orchestration. Cloud migration adds further complexity, as organizations must choose between self-managed clusters, cloud-native Kafka services, or hybrid deployments. Having explored the challenges, the natural next step understands how Kafka-driven architectures contribute to the future evolution of supply chain ecosystems. Hence, the next section provides a forward-looking perspective on technological advancements and innovation trends.

FUTURE DIRECTIONS: INTELLIGENT, AUTONOMOUS, AND PREDICTIVE SUPPLY CHAIN ECOSYSTEMS

Following the discussion on challenges, it becomes essential to transition toward emerging trends that will shape the next generation of supply chain visibility systems. The convergence of Kafka with edge computing, digital twins, AI-driven automation, blockchain, and 5G connectivity will propel real-time visibility into new dimensions of intelligence and autonomy. Digital twins will leverage continuous Kafka streams to simulate supply chain operations in real-time, allowing organizations to evaluate numerous scenarios with unprecedented accuracy. AI models trained on streaming data will anticipate disruptions, recommend optimized responses, and potentially automate entire decision chains, creating autonomous logistics networks. The integration of Kafka with blockchain systems will enable immutable traceability across supplier tiers, improving compliance and countering fraud. Edge processing will reduce latencies by preprocessing IoT data closer to the source, while Kafka will act as the backbone connecting distributed edge nodes to centralized analytics. As supply chains continue to expand globally, 5G networks will generate massive high-frequency data streams from mobile and stationary devices. Kafka's scalability and high throughput will be essential for harnessing this data, supporting applications such as autonomous delivery vehicles, drones, and fully automated warehouses. These future directions set the stage for the concluding section, which synthesizes the insights presented throughout this article.

CONCLUSION

Real-time supply chain visibility has become a strategic imperative for modern enterprises navigating global uncertainties, rising customer expectations, and increasing operational complexities. Kafka-based data streaming architectures offer a transformative approach to capturing, integrating, and analyzing continuous event streams that reflect the real-world state of supply chain operations. By leveraging Kafka's scalability, durability, fault tolerance, and event-driven design, organizations can break down data silos, enhance responsiveness, and enable predictive intelligence across end-to-end supply networks. Throughout this article, Kafka has been shown to support a wide array of supply chain functions, including transportation monitoring, inventory synchronization, production tracking, and predictive decision-making. While challenges exist in terms of integration, governance, and scalability, they can be addressed through strategic planning and advanced architectural patterns. Looking ahead, Kafka will serve as a foundational technology for intelligent, autonomous, and highly resilient supply chain ecosystems.

REFERENCES

1. Somapa, S., Cools, M., & Dullaert, W. "Characterizing supply chain visibility–a literature review." *The International Journal of Logistics Management* 29.1 (2018): 308-339.
2. Machado, D. S. M., Moreira, D., Castro, H., & Barão, A. "Real-time logistic monitoring: System prototype overview." *Journal of Information Systems Engineering and Management* 6.2 (2021).
3. Koberg, E., & Longoni, A. "A systematic review of sustainable supply chain management in global supply chains." *Journal of cleaner production* 207 (2019): 1084-1098.
4. Isah, H., Abughofa, T., Mahfuz, S., Ajerla, D., Zulkernine, F., & Khan, S. "A survey of distributed data stream processing frameworks." *IEEE Access* 7 (2019): 154300-154316.
5. Bousdekis, A., Lepenioti, K., Apostolou, D., & Mentzas, G. "A review of data-driven decision-making methods for industry 4.0 maintenance applications." *Electronics* 10.7 (2021): 828.
6. Chardonnens, T., Cudre-Mauroux, P., Grund, M., & Perroud, B. "Big data analytics on high Velocity streams: A case study." *2013 IEEE International Conference on Big Data*. IEEE, (2013).
7. Michelson, B. M. "Event-driven architecture overview." *Patricia Seybold Group* 2.12 (2006): 10-1571.
8. Tavana, M., Shaabani, A., Raeesi Vanani, I., & Kumar Gangadhari, R. "A review of digital

Publisher: SARC Publisher

transformation on supply chain process management using text mining." *Processes* 10.5 (2022): 842.
9. Ameer, T., & Valilai, O. F. "Cloud-native causal AI for supply chain KPI monitoring: A GCP framework to diagnose out-of-stock events." *Machine Learning with Applications* (2025): 100765.
10. Saleh, H. H., Chyad, A. K., Barakat, M., & Naamo, G. S. "Enhancing Business Operations Efficiency Thorough Predictive Analytics." *Journal of Ecohumanism* 3.5 (2024): 700-714.
11. Islam, A. M. M. S. "Enhancing supply chain resilience in Bangladesh through multimodal transport integration." *Decision Making and Analysis* (2025): 72-85.
12. Tian, X., Han, R., Wang, L., Lu, G., & Zhan, J. "Latency critical big data computing in finance." *The Journal of Finance and Data Science* 1.1 (2015): 33-41.
13. Sim, C., Zhang, H., & Chang, M. L. "Improving end-to-end traceability and pharma supply chain resilience using blockchain." *Blockchain in Healthcare Today* 5 (2022): 10-30953.
14. Balakrishnan, M., Malkhi, D., Wobber, T., Wu, M., Prabhakaran, V., Wei, M., & Zuck, A. "Tango: Distributed data structures over a shared log." *Proceedings of the twenty-fourth ACM symposium on operating systems principles*. (2013).
15. Carzaniga, A., Rosenblum, D. S., & Wolf, A. L. "Achieving scalability and expressiveness in an internet-scale event notification service." *Proceedings of the nineteenth annual ACM symposium on Principles of distributed computing*. (2000).
16. Cristian, F. "Understanding fault-tolerant distributed systems." *Communications of the ACM* 34.2 (1991): 56-78.
17. Javed, M. H., Lu, X., & Panda, D. K. "Characterization of big data stream processing pipeline: A case study using flink and kafka." *Proceedings of the fourth IEEE/ACM international conference on big data computing, applications and technologies*. (2017).
18. Bruckner, R. M., List, B., & Schiefer, J. "Striving towards near real-time data integration for data warehouses." *International Conference on Data Warehousing and Knowledge Discovery*. Berlin, Heidelberg: Springer Berlin Heidelberg, (2002).
19. Majumder, A., Shrivastava, N., Rastogi, R., & Srinivasan, A. “Scalable content-based routing in pub/sub systems.” *In IEEE INFOCOM* . (2009): 567-575). IEEE.
20. Elsaleh, T., Enshaeifar, S., Rezvani, R., Acton, S. T., Janeiko, V., & Bermudez-Edo, M. "IoT-Stream: A lightweight ontology for internet of things data streams and its use with data analytics and event detection services." *Sensors* 20.4 (2020): 953.
21. Wenzel, H., Smit, D., & Sardesai, S. "A literature review on machine learning in supply chain management." *Artificial Intelligence and Digital Transformation in Supply Chain Management: Innovative Approaches for Supply Chains. Proceedings of the Hamburg International Conference of Logistics (HICL), Vol. 27*. Berlin: epubli GmbH, (2019).
22. Lekić, M., Rogić, K., Boldizsár, A., Zöldy, M., & Török, Á. "Big data in logistics." *Periodica Polytechnica Transportation Engineering* 49.1 (2021): 60-65.
23. Hahn, C. K., Duplaga, E. A., & Hartley, J. L. "Supply-chain synchronization: lessons from Hyundai Motor Company." *Interfaces* 30.4 (2000): 32-45.
24. Van der Aa, H., Leopold, H., Mannhardt, F., & Reijers, H. A. “On the fragmentation of process information: Challenges, solutions, and outlook.” In *International Workshop on Business Process Modeling, Development and Support* (2015) 3-18). *Cham: Springer International Publishing*.
25. Boulos, P. F., Altman, T., Jarrige, P. A., & Collevati, F. "An event-driven method for modelling contaminant propagation in water networks." *Applied mathematical modelling* 18.2 (1994): 84-92.

Source of support: Nil; **Conflict of interest:** Nil.

Cite this article as:
Agarwal, R. “Real-Time Supply Chain Visibility Using Kafka-Based Data Streaming Architectures." *Sarcouncil Journal of Applied Sciences* 6.4 (2026): pp 39-45.

Publisher: SARC Publisher