




Edge-to-Cloud Data Engineering for Retail Inventory

Madhu Sathiri

Independent Researcher

sathirimac@gmail.com

Abstract

Inventory tracking is crucial for effective retail operations that rely on physical stocks. However, tracking is not simply about maintaining a current state of product stocks but also involves the management of transactions and the analysis of the past. Hence, the data engineering needs for inventory tracking require a distributed architecture that groups functionality into a set of data processing sub-systems connected through a set of shared data flows. Distributed data engineering architectures are examined in terms of all-in-one data streaming and lakehouse designs as well as more modular event-driven designs based on typical data streams for enterprise applications and data lakehouse semantics. Further discussion concerns product and stock data models as well as the elements needed for channel management in a sales system.

Retailers need to take many measures to ensure smooth shopping experiences for their customers. Among these critical actions, being able to manually or automatically monitor the state of physical stocks in a store is important to satisfy customer visits. The concept of inventory tracking is linked to this need; however, inventory tracking is not simply about maintaining the current state of product stocks. Rather, it includes management of sold and purchased stocks, background processes that update the stocks, and supporting transactional data. Being able to evaluate whether a certain product was made available for sale in time and under what conditions during its exposure time is crucial information for taking stock and sales channel decisions. Hence, the data engineering needs for inventory tracking require a distributed architecture that groups functionality into a set of data processing sub-systems connected through a set of shared data flows.

Keywords : Smart retail, inventory tracking, edge computing, Internet of Things, distributed data engineering, data streaming, event-driven architecture, lakehouse, lake architecture.

1. Introduction

Smart retail uses advanced technologies to create a seamless shopping experience for consumers. Real-time inventory tracking supports many of these technologies by detecting customers' shopping activities from entering the store to completing their purchases. Passive detection techniques based on radio frequency identification (RFID), computer vision, and sensor fusion substantially reduce the burden of inventory management. Data generated by these passive detection systems can be complex, heterogeneous, and produced at a large scale. The architecture for distributed data engineering must be carefully designed to ensure that the supply chain can be supported by data pipelines that combine data streaming with batch processing, handle heterogeneous data types, ensure data quality, manage data consistency, implement storage tiering, and integrate data into a unified model.

Data streaming and event-driven architectures enable support for disseminating detected customer activities to interested parties via a publish-and-subscribe mechanism. A data lakehouse architecture promotes the decoupling of storage and processing for stateful systems such as inventory tracking. A lake architecture augments data lakehouse concepts to enable keeping the latest stock-take and product information in an easily accessible format while using the lakehouse part for long-term storage and archival or analytical purposes. Specific data models expose the semantics of the product, stock, and transaction entities concerning their temporal and spatial dimensions.

Data streaming and event-driven architectures play a crucial role in modern data platforms by enabling real-time communication of detected customer activities to multiple interested systems through a publish-and-subscribe mechanism. In this approach, events such as purchases, product views, or inventory updates are streamed continuously and published to data channels where subscribed services—such as recommendation engines, analytics platforms, or notification systems—can react immediately. Complementing this, a data lakehouse architecture supports the decoupling of storage and processing, allowing stateful systems like inventory tracking to manage large volumes of operational data efficiently. Building further on this concept, a lake architecture enhances the lakehouse by maintaining the most recent stock-take and product information in a readily accessible format for operational use, while the lakehouse layer manages long-term storage, historical records, and analytical workloads. To ensure meaningful interpretation of the data, specialized data models define the semantics of product, stock, and transaction entities, incorporating both temporal dimensions (such as time-based changes in inventory or sales) and spatial dimensions (such as store or warehouse locations).





Fig 1: Real Time Data Architecture In Retail



1.1. Background and Significance

Increasingly automated retail stores continuously monitor the presence, availability, and use of products by consumers. Products are equipped with electronic identifiers, such as RFID tags, which are tracked by readers, either fixed at specific locations or carried by customers enabling more granular tracking of usage over time. Reading operations automatically update a central system, which maintains the presence and stock level of products in each store, as well as records sales transactions. These systems generate a continuous stream of events, which can be leveraged by analyzing instant changes of the stock level and by detecting out-of-stock situations.

Such systems operate a data hub that supports multiple analytical and operational business functions in near real-time. These functions include near-real-time stock-out detection enabling automatic replenishment management; detection of stock-outs on frequently purchased items; support of marketing campaigns aiming to boost sales of specific brands; transit monitoring of promotion items misplaced in other non-target stores; and hourly reporting of detailed sales transactions. Business rules capture the resilience of retail business processes and guide the development of dedicated pipelines.

Equation 1: Core inventory balance equation

1. Step 1: Start with stock at the previous observation time $t_{\{k-1\}}$. Denote it by $S_{\{p,l\}}(t_{\{k-1\}})$.
2. Step 2: Add all positive movements during the interval $(t_{\{k-1\}}, t_k]$, namely inbound replenishments $I_{\{p,l\}}(t_k)$ and valid returns $R_{\{p,l\}}(t_k)$.
3. Step 3: Subtract all negative movements during the same interval, namely sales/removals $O_{\{p,l\}}(t_k)$, damage/shrinkage $D_{\{p,l\}}(t_k)$, and any unresolved unknown adjustment $U_{\{p,l\}}(t_k)$.
4. Step 4: Combine the terms to obtain the discrete-time balance:

$$S_{\{p,l\}}(t_k) = S_{\{p,l\}}(t_{\{k-1\}}) + I_{\{p,l\}}(t_k) + R_{\{p,l\}}(t_k) - O_{\{p,l\}}(t_k) - D_{\{p,l\}}(t_k) - U_{\{p,l\}}(t_k)$$

2. Problem Statement and Requirements

Accurate and up-to-date knowledge about product stocks is essential for a retailer's business success. Unfortunately, retail inventories are often poorly maintained and recorded with significant delay and inaccuracies. A distributed data engineering architecture is proposed for smart retail inventory tracking that captures information about stocks and stock changes as they happen by means of automatic identification technologies. Such data are ingested, stored, and made available in a timely manner supporting a variety of use cases ranging from simple dashboards to AI models for demand forecasting. Real-time, near real-time, and batch processing architectures are employed, together with state, event, and temporal data models.

Data is captured by automatic identification technologies such as RFID, camera vision, and sensor networks. An RFID installation on a shelf exploits tag reads to identify products being taken from and placed on the shelf. Shelf capacity is enabled by a sensor network that detects when products are near the shelf's end. Additional tag reads from products in the vicinity of the register or scanned by customers contribute to the detection of sales transactions, while a camera vision system is installed behind the cashiers to detect when products are leaving and to generate re-stocking requests. Stock changes are ingested in real time, enriched using event-driven techniques, and stored in a real-time system for dashboarding and other use cases that operate on

known or near-actual states. Inventory states are used to answer "what is currently in stock" dashboards and also provide the expected answer for questions such as "what will be in stock in the next hour" if history repeats.





Fig 2: Problem Statement and Requirements of Distributed Data Engineering

2.1. Research design

This research uses an interpretive technology design approach to develop a distributed Data Engineering architecture tailored to enable smart retail inventory tracking. Three architectural qualities—real-time data streaming, unstructured or semi-structured data capture, and data quality management—are considered both singly and in combination and their implications examined through two semi-structured case studies: one exploring food loss in retail and the other focusing on product availability on grocery shelves. For both cases, managed physically distributed, logical cluster architectures combine heterogeneous mutable-data stores with shared-nothing processing. The solutions are deployed via clusters of widely available compute nodes, minimizing operational expenditures. Further, the system responsible for maintaining real-time product-state and -placement information applies a micro-batching technique that captures discrete events to mitigate inconsistencies due to unseen product-state changes. Finally, the architectures of the deployments are documented as precedents for other investigators.

Tracking product stock on supermarket shelves is important for consumer satisfaction and impacts purchasing decisions. Lack of availability on grocery shelves can lead to consumer dissatisfaction, impacting loyalty and long-term revenue. Tracking product stock and availability is essential for consumer satisfaction. A wide variety of items in food retail makes accurate stock monitoring a problem. Important product category for many retailers. The impact of food waste in retail is significant and was estimated at about 41 billion euros for the European retail sector in 2021.



Equation 2: Continuous-time form

Step 1: Suppose stock changes continuously because streaming events arrive at fine time granularity.

Step 2: Replace finite interval changes by rates $i_{\{p,l\}}(t)$, $r_{\{p,l\}}(t)$, $o_{\{p,l\}}(t)$, $d_{\{p,l\}}(t)$, and $u_{\{p,l\}}(t)$.

Step 3: The net rate of stock change is the sum of positive rates minus the sum of negative rates:

$$dS_{\{p,l\}}(t)/dt = i_{\{p,l\}}(t) + r_{\{p,l\}}(t) - o_{\{p,l\}}(t) - d_{\{p,l\}}(t) - u_{\{p,l\}}(t)$$

Step 4: Integrate both sides from time t_0 to t to recover the state equation:

$$S_{\{p,l\}}(t) = S_{\{p,l\}}(t_0) + \text{integral from } t_0 \text{ to } t \text{ of } (i+r-o-d-u) dt$$

3. Architectural Paradigms for Distributed Data Engineering

Data Engineering is a foundational pillar of Data Science that focuses on the acquisition, preparation, and support of data for analysis. Similar to Data Science in the wider sense, Data Engineering combines expertise and technologies from numerous disciplines and areas. Among these, Data Streams, Event-Driven Architecture, Data Lakes and Lakehouses are increasingly becoming of paramount importance for both surrounding developments in other domains and also within Data Engineering itself, moving the focus away from traditional batch-oriented pipeline processes. As a specific example, within the complex ecosystem of smart retail inventory-tracking products and systems, ultra-fast stream ingestion, semantic data integration, and wide-area test-data generation have emerged as the main areas of interest.

The Event-Driven Architecture layer covers everything from the generation of machine-generated data from the S-ETL store for operational processing in U-ETL ingest pipelines, to the subsequent real-time triggering of multiple stacks in Data Science and Business Intelligence. Both Data Lake and Lakehouse concepts unite the two extremes, fast-continual ingestion and ultra-slow reading, into a single, semantically structured entirety that supports both real-time event-driven applications and long-term Data Warehouse-like analysis.





Fig 3: High-level architecture design diagram for Inventory Management

3.1. Data Streaming and Event-Driven Architectures

Smart retail inventory tracking involves supporting applications with strict requirements on the freshness and completeness of inventory data. A possible violation of these requirements can be detected through the modelling of relevant concepts. Distributed data engineering paradigms that are effective in addressing these requirements are discussed.

Modern retail chains frequently adopt a smart retail approach, involving a combination of data-driven services and automations concentrated on customer engagement. Examples of services are chatbots and digital signage that provide visual and textual information on products. Critical for these services is the real-time availability of data describing products and stock levels. Such data are built and maintained by distributed data engineering systems sourced, for example, from inventory transactions and from Internet of Things (IoT) devices that track inventory movements outside of the physical chain, such as self-service checkouts, mobile applications, RFID sensors, or IoT atom boxes. However, these systems do not guarantee that the data are either complete or sufficiently fresh before being accessed by the services. In the case of stock levels, non-temporally consistent data can lead customers to perceive an out-of-stock condition when the stock level is in fact greater than zero. Stale data can substantially



degrade the customer experience. Stale information can affect the automatic reordering of products that run low, as well as the sales of products that will expire before being sold.

Equation 3: Event-driven micro-batch update

Step 1: The article mentions micro-batching to reduce inconsistencies from unseen product-state changes.

Step 2: Let a batch B_m contain events $e_1, e_2, \dots, e_n$ captured in a small window Δt_m .

Step 3: Let each event contribute a signed stock change $\Delta s(e_j)$.

Step 4: The batch contribution equals the sum of event deltas:

$$\Delta S(B_m) = \text{sum over } e_j \text{ in } B_m \text{ of } \Delta s(e_j)$$

Step 5: Apply that batch delta to the previous state:

$$S^{\{(m)\}}_{\{p,l\}} = S^{\{(m-1)\}}_{\{p,l\}} + \Delta S(B_m)$$

3.2. Data Lakehouse and Lake Architectures

The lakehouse architectural style supports hybrid transactional—analytical processing (HTAP) workloads that demand low-latency transactions while ensuring data correctness and quality. Lakehouse systems combine the benefits of data warehouses—supporting data quality, data consistency, efficient querying, and support for ACID transactions—with the benefits of data lakes—supporting semi-structured, unstructured, and low-cost storage services at scale.

The lake architecture supports data ingestion and processing in its undifferentiated state as much as possible to minimize transformation and preparation costs. However, unlike the lakehouse, the lake architecture does not support querying and processing with higher quality or correctness. For the inventory tracking system, the lake destination tier is primarily used for historical transactional data for analytics. The lake destination tier does not share the Web interface of the lakehouse; instead, its data is refreshed in batch mode during non-working hours. Lake-sensitive data, for example, product images require storing in the lakehouse for making decisions in the minute-level timescale ETL.

4. Data Models and Semantics for Inventory Tracking

The ability to achieve and sustain consistency in the real-time state of inventory is fundamentally determined by the semantics and experiments of the product, stock, and transactional entities used in the inventory tracking processes and in the ingestion, storage, and processing of the associated data. Elements of the state of inventory that are tracked explicitly—the entities stored, amount and location of stock in each location category, amount and location of on-the-move stock, locations where inventory transactions are to be processed presently, amount of stock sold, returned, and damaged in each location category during the transactions being processed at present—are naturally stored in the freshest of the data tiers. Any cross-sectional dimensions in the inventory state data and state data quality can change at any point in time. Elements that affect the accuracy of the real-time state of inventory and that cannot be visually detected or tracked presently, e.g., stock thefts and damages, are maintained as unknowns to the real-time tracking processes.

In addition to the semantic foundations that control data quality, the maintenance of quality and the validation of data on-the-fly before its usage in feeds to real-time operations and services (e.g., quality-driven short-term embeddings of the distributions of unknown elements) are critical to ensuring a correct and useful state of inventory at all times. Management of all detected data class and feed-source unknowns is essential for sustaining a consistent real-time state of stock in all categories, for valid estimation of stock in known-to-be-empty service categories (mainly expensive stock), and for the automation of fraud detection. Quality checks at the detailed data source level (e.g., condition checks on images of returned stock) can avoid harmful servicing of automated via more-mature-quality validation strategies. Validation checks at scales larger than those currently thought feasible, including any one of additive, error-controlled summation, and posterior distribution validation checks, can enable correct real-time servicing at desired scale without quality bottlenecks.




Fig 4: Real time data is important in Retail

4.1. Product, Stock, and Transactional Entities

Referring to product, stock, and transactional entities, an inventory-tracking datacube supporting stakeholder operations and business analytics considers products in a given set with identifiable properties and for which a transaction log exists. The product entity is expanded on with product class and attributes such as name, description, full price, discount percentage, barcode, image, e-commerce information, geographic provenance, and so on; product spatial-temporal provenance may also be specified to support stock replenishment logistics.

Inventory stock and associated transaction logs are generated by incoming and outgoing transactions affecting stock. Stock is represented by a state point table denoting at real time the quantity available at all locations across the product space-time. Inventory entries provide an origin-date product-quantity store-movement description, while outgoing transactions log a user-specified instant-based function transformation over a limited space-time region producing distinct and named—collectible or



delivered—stock-group items. A total transaction log tracks all incoming and outgoing transactions at all locations with user reference/debug data; at production origin, it traces produce origin-moving geography and advertisement.

Equation 4: Confidence-weighted inventory under uncertainty

Step 1: Suppose source i reports quantity q_i for a product-location pair.

Step 2: Assign confidence c_i in $[0,1]$ to that observation, where 1 means fully trusted and 0 means unusable.

Step 3: The effective contribution of source i is $c_i q_i$.

Step 4: Summing trusted contributions over all sources gives the confidence-weighted expected stock:

$$\hat{S}_{\{p,l\}}(t) = \sum_i c_i(t) q_i(t)$$

Step 5: If total confidence weights are to be normalized across heterogeneous sources, use the weighted average form:

$$\hat{S}_{\{p,l\}}(t) = (\sum_i c_i q_i) / (\sum_i c_i)$$


4.2. Temporal and Spatial Dimensions

Un magasin est un espace dans lequel des produits individuels peuvent être ajoutés et retirés par le biais d'une ou plusieurs transactions. Par conséquent, la structure de données Product doit également contenir une référence à l'espace d'un magasin et la structure de données Time doit contenir un horodatage de transaction. L'ajout d'une dimension temporelle à un magasin permet de mieux s'ajuster à un stock de produits qui ne sont pas inaltérables et qui sont soumis à des variations et à des ruptures de stock. Les données de stock à l'intérieur d'un magasin changent car des produits sont échangés et un magasin se vide ou se remplit au cours de l'année. Un bureau de tabac peut être à court de certains articles pendant trois jours, puis en recevoir de nouveau, ou avoir trop de certains produits, par exemple de la boisson gazeuse pendant la période de Noël, mais ne plus en avoir pendant l'été. Ces échelles de temps doivent également être intégrées à la structure des données. Les ruptures de stock peuvent également survenir pour une durée prolongée, comme pour les rayons qui ne sont pas souvent réassortis, tels que les jouets, les cosmétiques ou le matériel de pêche.

Pour un magasin donné, le produit et l'espace d'horloge représentent un groupe identifiant le stock. Par conséquent, le stock peut être décrit par une structure de données distincte associée à un magasin, à un produit et à un horodatage, avec les quantités nouvelles et supprimées comme propriétés. Un horodatage est nécessaire afin de tenir compte des périodes de rupture de stock.

5. Data Ingestion, Quality, and Consistency

Data ingestion pipelines acquire raw messages from sensors or inventory management systems. Automatic connectors handle well-defined streams and queues, as for RFID readers capturing product arrivals at the warehouse gate. Semantic enrichment fills in missing semantics for stores or items, possibly including bar code interpretation. Specialized tools monitor data quality and integrity in key domains. Messages for which entity definitions and semantics are unknown pop up as alerts requiring human attention.



Maintaining a consistent image of the inventory is complicated by asynchronously arriving messages of different types. Hardware arrivals update only the product volume in a store; reader message counts offer a stock snapshot only for the past hour; sales and removals are also important. Yet real-time updates are essential for triggering replenishment orders and alerting shop staff to theft. With very short-lived transactions, sale or removal consistency is relaxed: displayed stock is still the key. Taking these requirements into account, a temporary solution draws stock and sales data into a common NoSQL database dedicated to monitoring, augmented with time dimensions to validate product availability in different stores.

Equation 5: Inventory freshness and latency metrics

Step 1: For event e , define ingestion latency as the difference between processing time and event occurrence time.

$$L_e = t_{\text{process}}(e) - t_{\text{event}}(e)$$

Step 2: Over N events, average latency is:

$$L_{\text{avg}} = (1/N) \sum_{e=1}^N L_e$$

Step 3: Convert latency into a freshness score that decays as delay grows. A common exponential form is:

$$F_e = \exp(-\lambda L_e)$$

Step 4: Here $\lambda > 0$ controls how harshly delays are penalized. $F_e = 1$ means perfectly fresh; lower values mean staler data.

5.1. Ingestion Pipelines and Connectors

Data ingestion associates incoming data with appropriate syntactical, structural, and metadata specifications. Syntactical specifications define how input data is represented; structural specifications define how input data is organized; metadata specifications provide context to help validate associated data. Data ingestion solutions incorporate various technology components responsible for collecting and preparing raw data for potential analysis or action. Distributed data engineering architecture designs include ingestion pipelines enabling data collection and dissemination.

Data pipelines extract input from various connectors. Connectors can read from and/or write to data sources provided by other systems via APIs. Connecting software enables interaction with an external system's remote procedure call interface, receiving and/or sending data as problem-specific requests and responses. Stream-based information readily generated and published by a source can also be accessed via message-brokering software products, for the latter effectively facilitating the implementation of data streaming architectures. Data connectors capable of receiving and writing to streams and batches enabled by Apache Kafka, Apache NiFi, Amazon Kinesis, and Google Cloud's Pub/Sub are readily available.



5.2. Data Quality, Validation, and Unknowns


Despite the distributed and decentralised nature of the entire architecture, the implicit assumption remains that the state of the inventory at a point in time is consistent and valid, that an earlier known state is the same as the actually existing stock in the supermarket, and that transactions modify the stock in a deterministic way. Thus, no support for data validation in the ingestion pipelines has been specified. Yet in a real deployment, such an assumption would be unrealistic. Apart from traditional data validation mechanisms, another key aspect associated with the system design would be able to handle that part of the inventory management process and incoming data to which the necessary amount of trust and believe cannot be assigned.

The presence of uncertainty and unknowns in the inventory data is not an uncommon situation, due either to confidence on the existence of some individual or knowledge about its state being lower than the required threshold. Therefore, instead of keeping either the last known real state or just discarding confidence knowledge about the stock, a process based on the Bayesian Formalism can be applied to the stock, discounting each stock quantity according the associated confidence. With additional processing, also knowledge about singleton products being sold out in a given period of time or being refilled could be kept, enabling future usage of such information in the inventory process.

6. Storage Tiering and Metadata Management


The architectural design delineates two-tiered storage for inventory tracking data. The real-time storage tier of the lakehouse maintains the current stock state in a key-value format, facilitating efficient querying during operations such as smart shelf refills or customer purchases. Each read interface operates on a hot partition holding an up-to-date copy of the stock state, minimizing response times. Read operations leveraging in-memory key-value stores can also be deployed. Yet, the inventory tracking system frequently serves as a data source rather than a center of gravity, supplying real-time data to data warehouse or streaming solutions for customer experience enhancement.

The historical storage tier forms the data lake component. It stores bulk batch updates of the item's state and transaction-based batch updates of incoming and outgoing stock movements—crucial for buyers interested in sales trends. The state management for this component accommodates all data quality issues. Retrofit automates the joining and augmenting of data from source tables to feed the state change table. Ingestion into the state-change table can, indeed, bypass all conformance and transformation steps. Missing data and unknowns are identified, thereby ensuring data-validation steps are undertaken for the real-time copy primarily used by customer-facing applications.



6.1. Real-Time Storage for Inventory State

A distributed data engineering architecture can support inventory-tracking use cases based on a large mixed-city retail product collection—more than 1 million Sellable Product Items belonging to different retailer companies—and including customers, market baskets, transactions, product stock in different products of a grocery store, market baskets of a pharmacy, data of a woman’s shoe and handbag.



Although the used technologies support dynamic schema and also fingerprint an entity for flexible handling of unknowns, companies demand a high operational certainty to improve customer experience. Data on Supply Chain changes less rapidly than hourly, so the solution uses a real-time data-store and entry to maintain the live inventory state and a data lake to create a better self-service accessible customer shopping experience, enabling Visit Data Science, addressing the global-trend of Improve Experience with AI. An Event-Driven Architecture uses an Event Source pattern for initial dynamic data ingestion, dynamically storing the inventory product information inside a Lake Storage and maintaining their integrity using a Data-Transformation Service. For out-of-grocery-store and fixed-product Fails forms, a Data-Quality pattern applies automated validation, logic entries in the Data-Quality area of the Supply Chain Configuration Management Tool. Data consistency and integrity are also ensured in standard deployment and handling of unknowns for the transient changes in other Businesses. Popular Deployment cleaning management also provides image sourcing in brand-controlled environments and service on-sell information for audio.

Real-time ingestion and storage of Inventory Supply Chain status enable a next-step Connected-Intelligence use case: providing an Inventory Clothing Management Tool to support company Aprize Clothing. These Supply Chain data are massive and cannot provide spatial competence to Shopping Intelligence use cases.


6.2. Historical Data for Analytics

The architecture of retail inventory tracking systems is typically segregated into distinct storage tiers with varied technology implementations based on different service level requirements. A retained real-time view of all active stock units enables near instantaneous determination of available inventory position during a customer shopping session. While it must be served by high-frequency and performant technology (e.g., NoSQL in memory), tracking systems also tend to accumulate large volumes of historical data that preserve a detailed record of multiple years of transactions spanning every product attribute. This analytical tail of fine-grained transaction history is usually built up as a combination of batch loads and historicized delta synchronizations, consistently with the broader Data Lakehouse paradigm. However, the high ingestion frequency and transactional nature of the data can lead to additional challenges around volatility, skew, and cold-path resource sourcing.

Data architects are increasingly seizing similar opportunities to capitalize on the Data Cloud paradigm, leveraging externalized storage and compute functionality to elastically adapt data volumes to varying time-lengths of reversible cold-path and volatile write workloads. Such capability improves economics, with strategic target-and-park placement of the least expensive and performance-sensitive materialized views enabled by remote storage over the data. Unionized regular views per product and stock unit type allow efficient data discovery and automatic query steering to the most performant resources.

7. Conclusion

Emerging trends in distributed data engineering solutions that unify the analysis of inventory states and transactional data were reviewed. Smart retail environments require inventory routing through shop floors that remain connected in a spatial-temporal stream within a product-stock semantics. The integration of inventory and sales data is crucial for keeping stock levels consistent with temporal patterns in sales data. Processing constraints may arise due to the data- ingestion rate, requiring a schema-less architecture supporting diverse cash register terminals. Recent advances in the streaming-batch continuum and data-lake architecture allow these challenges to be addressed.



The increasing integration of data- engineering solutions and data-science activities toward a Shared-Real-Time layer has the potential to transform the application of analytics in various domains. Although the generation of prediction models continues to be a stimulus for data-science activity, novel data-engineering approaches that support the analytical consumption in real time move to the core of these business intelligence environments. They also assume a foundational role in reinforcement-learning practices that govern control processes.


Equation	Purpose	Link to article concept
$S_k = S_{\{k-1\}} + I + R - O - D - U$	Real-time stock state update	Transactions plus current inventory state
$dS/dt = i + r - o - d - u$	Continuous stream formulation	Streaming/event-driven architecture
$\Delta S(B_m) = \sum \Delta s(e)$	Micro-batch update	Micro-batching for consistency
$\hat{S} = \sum c_i q_i$	Confidence-weighted stock	Bayesian handling of unknowns
$L_e = t_{process} - t_{event}$	Latency measure	Freshness requirement
$C = N_{received} / N_{expected}$	Completeness measure	Data quality monitoring
$ROP = \mu_{\tau} + z\sigma_{\tau}$	Replenishment trigger	Automatic replenishment
$P(\text{stockout}) = 1 - \sum e^{(-\lambda T)} (\lambda T)^k / k!$	Risk of running out	Availability forecasting

Table : Compact summary tables

7.1. Emerging Trends


The inventory tracking problem for retail environments is well-suited for experimental deployment of smart retail solutions. During the last few years, several new paradigms and systems for distributed streaming data engineering have appeared. Ready-to-use technologies, including ready-to-use SLAs for various types of products, available connectors or ingestion pipelines, and world sales data feeding a data lake, support the experiments and promote conditions to upgrade a project-proof architecture to a state-of-the-art distributed streaming architecture.

Emerging technologies promise massive adoption in production environments. Data streaming and event-driven architectures enable elastic real-time processing of sudden workloads that are unpredictable in volume and frequency. Lakehouse architectures combine the low-cost storage tiering of data lakes with the performance and schema evolution capabilities of data warehouses. Supported adoption may thus drive customers' frictionless consumption. With minimal specialisation, automated ingestion pipelines can connect supply-side sales stream products to their respective validation procedures and mark resultant violations as unknowns. Unknowns predict products for which sales flow suddenly deviate from common behaviour. Products consuming stale inventories spawn alerts based on temporal and spatial sales concentration shaped by external factors and managed in near real time through controlled false negatives.




8. References


1. Ali, S. M., Ashraf, A., Taqi, H. M. M., Ahmed, S., Rob, S. M. A., Kabir, G., & Paul, S. K. (2023). Drivers for Internet of Things (IoT) adoption in supply chains: Implications for sustainability in the post-pandemic era. *Computers & Industrial Engineering*, 183, 109515.
2. Aravind, R., & Surabhii, S. N. R. D. (2023). Harnessing Artificial Intelligence for Enhanced Vehicle Control and Diagnostics.
3. Ahmad, S., Shakeel, I., Mehfuz, S., & Ahmad, J. (2023). Deep learning models for cloud, edge, fog, and IoT computing paradigms: Survey, recent advances, and future directions. *Computer Science Review*, 49, 100568.
4. Badakhshan, E., & Ball, P. (2023). Applying digital twins for inventory and cash management in supply chains under physical and financial disruptions. *International Journal of Production Research*, 61(15), 5094–5116.
5. Inala, R. (2023). Revolutionizing Customer Master Data in Insurance Technology Platforms: An AI and MDM Architecture Perspective. *International Journal of Finance (IJFIN)-ABDC Journal Quality List*, 36(6), 579-606.
6. Barbuto, V., Savaglio, C., Chen, M., & Fortino, G. (2023). Disclosing edge intelligence: A systematic meta-survey. *Big Data and Cognitive Computing*, 7(1), 44.
7. Bhandal, R. (2022). The application of digital twin technology in operations and supply chain management: A bibliometric review. *Supply Chain Management: An International Journal*, 27(2), 182–206.
8. Dolgui, A., & Ivanov, D. (2022). 5G in digital supply chain and operations management: Fostering flexibility, end-to-end connectivity and real-time visibility through Internet-of-Everything. *International Journal of Production Research*, 60(2), 442–451.
9. Fatorachian, H., & Kazemi, H. (2021). Impact of Industry 4.0 on supply chain performance. *Production Planning & Control*, 32(1), 63–81.
10. Kolla, S. H., & Loganathan, R. (2023). Cloud-Native Deep Learning Architectures For Secure Generative AI Deployment In Enterprise Workflow Platforms. *Journal of International Crisis and Risk Communication Research*, 603-618.



11. Gebhardt, M., Kopyto, M., Birkel, H., & Hartmann, E. (2021). Industry 4.0 technologies as enablers of collaboration in circular supply chains: A systematic literature review. *International Journal of Production Research*.
12. Jan, M. A., Zakarya, M., Khan, M., Mastorakis, S., & others. (2021). An AI-enabled lightweight data fusion and load optimization approach for Internet of Things. *Future Generation Computer Systems*, 122, 40–51.
13. Aitha, A. R. (2021). Dev Ops Driven Digital Transformation: Accelerating Innovation In The Insurance Industry. *Journal of International Crisis and Risk Communication Research*.
14. Karakas, S., Acar, A. Z., & Kucukaltan, B. (2022). Blockchain adoption in logistics and supply chain: A literature review and research agenda. *International Journal of Production Research*, 60, 8193–8216.
15. Khan, Y., Su'ud, M. B. M., Alam, M. M., Ahmad, S. F., Ahmad, A. Y. A. B., & Khan, N. (2023). Application of Internet of Things (IoT) in sustainable supply chain management. *Sustainability*, 15(1), 694.
16. Reddy, V. A. R. (2022). Designing Fault-Tolerant Data Ingestion Pipelines for High-Volume Healthcare Transactions. *Frontiers in Health Informatics*, 11, 861-889.
17. Kong, L., Tan, J., Huang, J., Chen, G., Wang, S., Jin, X., Zeng, P., Khan, M., & Das, S. K. (2022). Edge-computing-driven Internet of Things: A survey. *ACM Computing Surveys*, 55(8), 1–41.
18. Laroui, M., Nour, B., Moun gla, H., Cherif, M. A., Afifi, H., & Guizani, M. (2021). Edge and fog computing for IoT: A survey on current research activities & future directions. *Computer Communications*, 180, 210–231.
19. Mangalampalli, B. M. Generative AI Applications In Healthcare Data Mart Design And Optimization.
20. Liu, Z., & Nishi, T. (2023). Data-driven evolutionary computation for service constrained inventory optimization in multi-echelon supply chains. *Complex & Intelligent Systems*, 10, 825–846.
21. Maheshwari, P., & Kamble, S. (2022). The application of supply chain digital twin to measure optimal inventory policy. *IFAC-PapersOnLine*, 55(10), 2324–2329.



22. Mendez, J., Bierzynski, K., Cuéllar, M. P., & Morales, D. P. (2022). Edge intelligence: Concepts, architectures, applications, and future directions. *ACM Transactions on Embedded Computing Systems*, 21(5), 1–41.
23. Mattaparthi, R. (2022). Engineering Predictive Industrial Systems Through IoT-Driven Asset Monitoring and Machine Learning Prognostics. *International Journal of Future Innovative Science and Technology (IJFIST)*, 5(1), 7790.
24. Ovezmyradov, B., & Kurata, H. (2022). Omnichannel fulfillment and item-level RFID tracking in fashion retailing. *Computers & Industrial Engineering*, 168, 108108.
25. Park, H., Choi, D. G., & Min, D. (2023). Adaptive inventory replenishment using structured reinforcement learning by exploiting a policy structure. *International Journal of Production Economics*, 266, 109029.
26. Aravind, R., Surabhi, S. N. D., & Shah, C. V. (2023). Remote Vehicle Access: Leveraging Cloud Infrastructure for Secure and Efficient OTA Updates with Advanced AI. *Eur. Econ. Lett*, 13, 1309-1319.
27. Perez, H. D., Wassick, J. M., & Grossmann, I. E. (2022). A digital twin framework for online optimization of supply chain business processes. *Computers & Chemical Engineering*, 166, 107972.
28. Pournader, M., Ghaderi, H., Hassanzadegan, A., & Fahimnia, B. (2021). Artificial intelligence applications in supply chain management. *International Journal of Production Economics*, 241, 108250.
29. Qader, G., Junaid, M., Abbas, Q., & Mubarik, M. S. (2022). Industry 4.0 enables supply chain resilience and supply chain performance. *Technological Forecasting and Social Change*, 185, 122026.
30. Aravind, R., Surabhi, S. N. D., & Shah, C. V. (2023). Remote Vehicle Access: Leveraging Cloud Infrastructure for Secure and Efficient OTA Updates with Advanced AI. *European Economic Letters (EEL)*, 13 (4), 1308–1319. Retrieved from <https://www.european-economics.com/>
31. Rosendo, D., Costan, A., Valdúriez, P., & Antoniu, G. (2022). Distributed intelligence on the Edge-to-Cloud Continuum: A systematic literature review. *Journal of Parallel and Distributed Computing*, 166, 71–94.



32. Seyedan, M., Mafakheri, F., & Wang, C. (2023). Order-up-to-level inventory optimization model using time-series demand forecasting with ensemble deep learning. *Supply Chain Analytics*, 3, 100024.
33. Yandamuri, U. S. (2023). An Intelligent Analytics Framework Combining Big Data and Machine Learning for Business Forecasting. *International Journal Of Finance*, 36(6), 682-706.
34. Taj, S., Imran, A. S., Kastrati, Z., Daudpota, S. M., & others. (2023). IoT-based supply chain management: A systematic literature review. *Internet of Things*, 24, 100982.
35. Wang, L., Deng, T., Shen, Z.-J. M., Hu, H., & Qi, Y. (2022). Digital twin-driven smart supply chain. *Frontiers of Engineering Management*, 9(1), 56–70.
36. Inala, R. (2023). Big Data Architectures for Modernizing Customer Master Systems in Group Insurance and Retirement Planning. *Educational Administration: Theory and Practice*, 29(4), 5493-5505.