

# The Physical Store as a Real-Time Agent Service Node

Senthil Kumar Muthu

Independent Researcher

**Abstract:** The rapid growth of agentic commerce, in which autonomous artificial intelligence agents execute retail transactions on behalf of human consumers, is creating a structural demand that physical retail stores are not architecturally prepared to meet. This paper's primary objective is to define an architectural framework that transforms physical retail stores into real-time agent service nodes capable of responding to machine-initiated inventory and fulfillment queries with the low latency and high reliability that agent-driven transactions require. The research design combines an architectural analysis of current radio-frequency identification (RFID) deployment patterns in retail store operations with a proposed reference architecture, the Store Inventory Service, that unifies four infrastructure capabilities behind a single application programming interface (API) layer: item-level RFID accuracy, location-aware inventory mapping, fulfillment channel routing intelligence, and real-time store capacity signaling. The methodology further introduces a Store Readiness Assessment Framework across five operational dimensions, providing retailers a diagnostic model for evaluating the gap between current store technology states and an agent-ready target state, with a measured metric, maturity levels and thresholds for each dimension. A prototype exercised against 250,000 simulated agent queries cuts false promises from 6.48% to zero and p95 latency from 1,079 ms to 53 ms against a cycle-count baseline. The main outcomes identify that achieving reliable agent-query responses requires high aggregate inventory accuracy, fine-grained item-location resolution, low-latency API response times suitable for machine-speed transaction flows, and a graceful degradation strategy for handling RFID infrastructure disruptions without excluding the store entirely from agent-directed fulfillment routing. The paper concludes that the physical store's competitive position in an agent-mediated commerce landscape depends directly on its ability to expose real-time, machine-readable inventory and fulfillment state through a well-designed API architecture, and that industry-wide interoperability standards will be necessary to prevent agent commerce readiness from becoming disproportionately concentrated among the largest, best-resourced retailers.

**Keywords:** Agentic commerce, RFID, real-time inventory, store readiness, omnichannel retail, API architecture

---

## 1. Introduction

The physical retail store occupies a paradoxical position in the emerging landscape of agentic commerce. On one hand, physical stores represent the most valuable fulfillment asset in omnichannel retail: they provide same-day availability, immediate pickup, and an in-person shopping experience no digital surface can replicate. On the other hand, physical stores are the least machine-readable component of the retail commerce estate, with inventory data managed by systems designed for human operators on human timescales. When an autonomous AI agent attempts to determine whether a specific item is available for same-day pickup, it encounters infrastructure that was never designed to answer that question at machine speed, with machine reliability, in a machine-readable format.

This mismatch will become increasingly significant as agentic commerce scales. Retailers whose physical stores can respond to agent queries in real time will capture agent-directed traffic that competitors with machine-unreadable stores cannot serve, regardless of actual inventory or service quality [14]. Emerging AI-driven retail systems already demonstrate the commercial value of tightly coupling real-time tracking with inventory forecasting at the point of interaction [8], reinforcing that the store layer, not just the digital storefront, must become a first-class participant in machine-initiated commerce. The transformation of the physical store into a real-time agent service node is an architectural imperative, not a future consideration, for omnichannel retailers competing in an agent-mediated



commerce environment. This paper proposes the Store Inventory Service, a reference architecture that exposes four infrastructure capabilities as a unified API layer, and a Store Readiness Assessment Framework for evaluating the transition from current store technology states to this agent-ready target.

## 2. Literature Review

RFID technology has been deployed in retail store operations for over two decades, evolving from a supply chain tracking tool into a store-level inventory management capability [3], [4]. In the early stages of its use, RFID technology was mainly employed in the logistics industry for identifying pallets and boxes during the reception of goods and transfers between facilities. This indicated a significant advancement in terms of tracking the supply chain, but the technology's uses were limited [1]. Reducing costs associated with RFID readers, along with the fact that the implementation of item level tags became commercially viable allowed for marking items which included wide range of possibilities such as cycle counting, loss prevention, and receiving accuracy verification [2], [5].

Despite these advances, a fundamental gap remains between current store technology capabilities and the requirements of agent-initiated commerce. This gap has four dimensions. First, most RFID deployments are optimized for periodic batch inventory sweeps rather than continuous updates, which is adequate for human-paced workflows but insufficient for agent queries requiring near-real-time confirmation [3]. Second, current RFID deployments typically provide zone-level location data rather than the precise location resolution required to direct an associate to retrieve a specific item [13]. Third, store fulfillment capabilities such as same-day pickup and ship-from-store are typically maintained in internal systems not exposed as queryable APIs [9], [10], [11]. Fourth, physical stores have no established mechanism for signaling their real-time operational capacity to external systems [15], [16], [17]. Together, these four gaps define the agent readiness problem for physical retail stores that this paper addresses.

## 3. Method / Algorithm

### 3.1 Four Infrastructure Capabilities for Agent-Ready Physical Stores

Transforming a physical retail store into a real-time agent service node requires four specific infrastructure capabilities. Item-Level Real-Time RFID Accuracy is the foundation: the key metric is Aggregate Inventory Accuracy, the degree to which an item's RFID-reported state matches its actual physical state. Sustained high accuracy is achievable with well-designed reader networks and regular calibration [4], [5]; below an acceptable threshold, false-positive and false-negative rates in agent query responses become commercially significant. Location-Aware Inventory Mapping answers the follow-up question of where in the store an item is located, requiring resolution fine enough to direct an associate to the item efficiently, achieved through fixed overhead reader triangulation or handheld reader-assisted location recording [13].

Fulfillment Channel Routing Intelligence exposes fulfillment eligibility and capacity as queryable service parameters, combining item-level fulfillment rules from the product catalog with store-level capacity state that changes throughout the day [9], [10]. Real-Time Store Capacity Signaling is the fourth and most forward-looking capability, communicating the store's overall capacity for receiving and processing agent-initiated orders, enabling intelligent load balancing across multiple stores in a geographic area [17].

### 3.2 Proposed Reference Architecture: The Store Inventory Service

The Store Inventory Service is a unified API layer exposing all four capabilities as a single, versioned, authenticated endpoint, following the general pattern of decomposing monolithic IoT platforms into independently deployable, API-first services [6]. The architecture is built on three principles: the Single Query Contract, in which an agent obtains all information needed for a fulfillment routing decision in one API call; the Low-Latency Response Guarantee, requiring a response fast enough to fit within an agent's overall transaction latency budget; and the Confidence Scoring Principle, in which every response includes a relative confidence indicator reflecting the recency and reliability of underlying data. The architecture consists of three layers: a Data Aggregation Layer continuously ingesting RFID reader events, point-of-sale transactions, and fulfillment system events; an Inventory State Store

optimized for point-query access with minimal query latency; and an API Layer exposing both the Inventory State Store and Fulfillment Routing Service with rate limiting, caching, and graceful degradation.

The RFID data pipeline ingests raw read events from fixed overhead readers, portal readers, and handheld readers, processing them through deduplication, confidence scoring, and location inference. Confidence scoring assigns a reliability weight based on read recency, signal strength, and cross-reader consistency, enabling agents to set different confidence thresholds for different use cases — a same-day pickup query might require a higher confidence threshold than a general availability check intended for product discovery. Cost and traceability considerations in multi-technology data pipelines that combine RFID with adjacent tracking mechanisms have been analyzed in supply chain contexts outside retail as well, underscoring that the integration overhead of a unified data pipeline is a general architectural concern rather than one unique to any single industry [7]. A store's real-time state can further be understood as a continuously updated digital twin of the physical retail environment, extending prior work on IoT-based store digital twins from a primarily operational tool into an agent-facing service layer [19].

**Table 1. Four Infrastructure Capabilities Required for Agent-Ready Physical Retail Stores**

| Capability                               | Description                                                                                                                                  | Requirement for Reliable Agent Queries                                | Agent Query Type Enabled                                                                  |
|------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------|-------------------------------------------------------------------------------------------|
| Item-Level Real-Time RFID Accuracy       | Each item is uniquely identifiable by its RFID tag, with presence, location zone, and availability queryable within a defined refresh window | Sustained high aggregate inventory accuracy; frequent refresh         | Availability confirmation: is this item currently in this store?                          |
| Location-Aware Inventory Mapping         | Item location recorded at zone or position level, enabling directed retrieval                                                                | Fine-grained location resolution; timely location-association updates | Directed fulfillment: where is this item so an associate can retrieve it?                 |
| Fulfillment Channel Routing Intelligence | Fulfillment capabilities exposed as queryable service parameters agents can evaluate per item and order                                      | Real-time capacity state; channel eligibility rules per item type     | Fulfillment routing: can this item be fulfilled via same-day pickup here?                 |
| Real-Time Store Capacity Signaling       | Store exposes current operational capacity for agent-initiated orders as a machine-readable signal                                           | Frequent capacity refresh; low-latency queue-depth signaling          | Load balancing: which nearby store has the best combination of availability and capacity? |

### 3.3 Prototype Implementation and Validation

To test whether the proposed architecture is feasible and whether data currency is the constraint that governs its usefulness, the three layers described above were implemented as a discrete-event prototype and exercised against simulated item-level RFID data. The prototype models a single store carrying 12,000 stock-keeping units with sales velocity drawn from a lognormal distribution, twelve trading hours per day, a 48-hour replenishment cadence, a 2% receiving-error probability and a 1.5% unobserved-removal rate, with fixed and portal readers operating at 97% read accuracy and 0.4% of tags systematically unread. Three data-currency architectures were instantiated behind the identical query contract: A1, in which the inventory record is reconfirmed only by periodic manual cycle counting; A2, in which an overnight RFID sweep reconfirms it every twelve hours; and A3, the Store Inventory Service with continuous ingest, which reconfirms it every five minutes. Because all three answer the same contract with the same physical process underneath, any difference in what an agent receives is attributable to data currency alone.

Table 2 gives the query contract the prototype implements, exposed as one versioned, authenticated route, GET /v1/stores/{store\_id}/items/{sku}, which returns the listed fields as a single JSON object; a batch form of the same route answers the multi-item comparison an agent makes when it evaluates a basket. One call scoped to a store and an item returns sellable quantity, location, eligible fulfillment channels, a confidence value, the time the state was last confirmed by a read event, and a flag marking answers served from last-confirmed state during a reader outage. The confidence value and the degraded flag are what allow an agent to apply its own risk policy without a second call, which is the practical expression of the Single Query Contract principle.

**Table 2. Store Inventory Service Query Contract Implemented by the Prototype**

| Contract Element    | Type      | Description                                                                                       | Example                 |
|---------------------|-----------|---------------------------------------------------------------------------------------------------|-------------------------|
| store_id            | string    | Store the query is scoped to                                                                      | S-0042                  |
| sku                 | string    | Item identifier being queried                                                                     | SKU-8391104             |
| quantity_available  | integer   | Units the record holds as sellable at query time                                                  | 3                       |
| location            | string    | Location resolution for the item, at the finest level the store supports                          | ZONE-W2 / FIX-14        |
| fulfillment_options | array     | Channels the store can honour for this item, with current capacity applied                        | pickup, ship_from_store |
| confidence          | float 0-1 | Reliability of the answer, weighted by read recency, signal strength and cross-reader consistency | 0.97                    |
| as_of               | timestamp | Time the item state was last confirmed by a read event                                            | 2026-03-04T15:22:07Z    |
| degraded            | boolean   | Set when the answer is served from last-confirmed state during a reader or network outage         | false                   |

Each architecture answered 250,000 independent availability queries. The harness holds the physical shelf state separately from the inventory record, so every answer is scored against ground truth. An answer was scored as a false promise when the service reported an item as available and the shelf did not hold it, which is the failure an agent transaction cannot recover from once a customer commitment has been made. Response latency was modelled per layer across aggregation, state lookup and API serialization, against a 500 ms budget for the store-query step of an agent transaction and a 1,500 ms hard timeout. Because an agent comparing several nearby stores issues those queries in parallel and waits for the slowest answer, fan-out compliance was measured as the share of five-store queries in which every response landed inside the budget. Table 3 reports the results.

**Table 3. Prototype Validation Results Across Three Data-Currency Architectures**

| Measure (250,000 queries per architecture) | A1 Periodic Cycle Count | A2 Overnight RFID Sweep | A3 Store Inventory Service |
|--------------------------------------------|-------------------------|-------------------------|----------------------------|
|                                            |                         |                         |                            |

|                                             |                |              |             |
|---------------------------------------------|----------------|--------------|-------------|
| Mean age of the inventory record            | 24 days        | 6 hours      | 2.5 minutes |
| Aggregate unit-level inventory accuracy     | 74.9%          | 97.5%        | 99.6%       |
| In-stock answer rate                        | 91.6%          | 95.9%        | 96.0%       |
| False-promise rate on availability answers  | 6.48%          | 0.05%        | 0.00%       |
| False-promise rate at one unit on hand      | 10.6%          | 0.10%        | 0.00%       |
| Query latency, p50 / p95                    | 491 / 1,079 ms | 100 / 200 ms | 36 / 53 ms  |
| Single query inside the 500 ms agent budget | 51.6%          | 99.97%       | 100.0%      |
| Five-store fan-out inside the 500 ms budget | 3.7%           | 99.8%        | 100.0%      |
| Query timeout rate at 1,500 ms              | 1.13%          | 0.00%        | 0.00%       |

The results locate the constraint precisely. Continuous ingest raises aggregate unit-level accuracy by 24.7 percentage points over periodic cycle counting and the in-stock answer rate by 4.5 points, but the consequential change is in the false-promise rate, which falls from 6.48% to zero across 250,000 queries; at a single unit on hand, the traditional architecture promises an item it cannot supply in 10.6% of answers. Latency separates the architectures by a further order of magnitude: p95 response falls from 1,079 ms to 53 ms, and the five-store parallel comparison an agent actually performs completes inside the 500 ms budget in 3.7% of cases under periodic counting against 100% under continuous ingest. Confidence scoring also turns out to be strong enough to serve as a routing policy in its own right: across a mixed fleet of 35% A1, 40% A2 and 25% A3 stores, declining to answer below a confidence of 0.95 still serves 91.7% of queries, and does so at 99.86% precision.

Two limits on this evidence should be stated plainly. The prototype exercises the architecture against simulated read streams and not against a live store estate, so the latency figures compare architectures under identical assumptions; they are not field measurements. The read-accuracy and shrinkage parameters are drawn from the published literature cited in Section 2, so absolute values will move with a retailer's own reader density and category mix. What holds across the range tested is the ordering: the false-promise rate climbs from zero to 6.78% as the reconfirmation interval is swept from five minutes to 575 hours, and rises at every step as the unobserved-removal rate is swept from 0.5% to 5%. The simulation is deterministic from a fixed random seed, and the harness and the full result set are available from the author on request.

## 4. Results and Discussion

### 4.1 Store Readiness Assessment Framework

Transforming an existing store into an agent-ready service node is an incremental, multi-year journey. The Store Readiness Assessment Framework evaluates a store's current state across five dimensions: Inventory Accuracy, Inventory Location Visibility, Fulfillment Channel Exposure, Store API Readiness, and Operational Technology Integration. For each dimension, the framework defines the current state typical of traditional stores, the target agent-ready state, and a key transition indicator — for example, moving from infrequent, cycle-count-based accuracy checks to continuous, RFID-driven accuracy monitoring updated in near-real time, an approach consistent with research showing that structured count-prioritization procedures materially improve inventory accuracy outcomes in retail store environments [18].

**Table 4. Store Readiness Assessment Framework — Traditional Store to Agent-Ready Store**

| Readiness Dimension                | Current State                         | Target State                              | Key Transition Indicator                                          |
|------------------------------------|---------------------------------------|-------------------------------------------|-------------------------------------------------------------------|
| Inventory Accuracy                 | Cycle-count-based; infrequent updates | RFID-driven; near-real-time updates       | Sustained high accuracy across item categories                    |
| Inventory Location Visibility      | Zone-level; manual or absent mapping  | Item-level, queryable via API             | Fine-grained location resolution with low update latency          |
| Fulfillment Channel Exposure       | Internal systems only, not queryable  | Real-time queryable service endpoints     | Low-latency fulfillment query response; frequent capacity refresh |
| Store API Readiness                | No machine-queryable APIs             | Standardized, authenticated API layer     | High uptime; low-latency query response                           |
| Operational Technology Integration | Independent systems, limited sharing  | Unified real-time inventory state service | Low end-to-end latency from transaction to queryable state update |

The dimensions in Table 4 become actionable once each carries a measured metric and threshold values, and Table 5 supplies them. Every dimension is scored on a five-level maturity scale, from Level 0, the state of a store with no agent-facing capability at all, to Level 4, the state at which that dimension no longer limits agent-directed fulfillment. Level scores are combined into a composite readiness score as a weighted mean of the five levels scaled to 100, with the weights in the table reflecting how directly each dimension determines whether an agent receives a correct answer. Four bands interpret the composite: 0 to 39 Not Ready, at which a store should not be exposed to agent traffic at all; 40 to 59 Foundational; 60 to 79 Transitional, at which the store can be exposed behind a confidence threshold of the kind described in Section 3.3; and 80 to 100 Agent-Ready. The thresholds are anchored on the measurements in Section 3.3: the agent-ready boundary in each dimension is an operating point the prototype reached, and the lower levels mark the states it was measured against.

**Table 5. Measurable Readiness Metrics, Maturity Levels and Threshold Values for Each Readiness Dimension**

| Readiness Dimension (Weight)        | Measured Metric                                                                                | Threshold Values by Maturity Level                                                                                                                                                                  |
|-------------------------------------|------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Inventory Accuracy (30%)            | Aggregate unit-level inventory accuracy, audited against a full physical count                 | L0 below 90%; L1 90 to 95%; L2 95 to 97.5%; L3 97.5 to 99%; L4 at or above 99.5%                                                                                                                    |
| Inventory Location Visibility (15%) | p90 associate search time to locate a queried item, against a 6-minute pick service level      | L0 no location data; L1 store-level only; L2 zone-level, p90 within 11 minutes; L3 fixture-level, p90 within 6 minutes; L4 item-level, p90 within 3 minutes and service-level misses at or below 1% |
| Fulfillment Channel Exposure (20%)  | False-promise rate on the availability answers used for fulfillment routing                    | L0 no queryable channel; L1 above 5%; L2 at or below 2%; L3 at or below 0.5%; L4 at or below 0.1%                                                                                                   |
| Store API Readiness (20%)           | p95 single-query latency, and share of five-store fan-out queries inside a 500 ms agent budget | L0 no API; L1 p95 above 1,000 ms; L2 p95 within 500 ms; L3 p95 within 200 ms and fan-out at or above 99%; L4 p95 within 100 ms and fan-out at 100%                                                  |

|                                             |                                                                                           |                                                                                                                                                                                                  |
|---------------------------------------------|-------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Operational Technology Integration<br>(15%) | Mean age of the inventory record,<br>expressed as the interval between<br>reconfirmations | L0 above 30 days; L1 within 30<br>days; L2 within 24 hours; L3<br>within 19 hours, the longest<br>interval that holds a 0.1% false-<br>promise budget; L4 within 5<br>minutes, continuous ingest |
|---------------------------------------------|-------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

#### 4.2 Latency Requirements and Graceful Degradation

The latency requirements for agent-facing store inventory queries are significantly more demanding than for human-facing queries. A human associate can tolerate a noticeably longer response without meaningful impact on their work; an AI agent operating within a defined multi-step transaction latency budget requires a much faster response, since the combined latency across all API calls in a single agent transaction must stay within a tight overall budget [20]. Graceful degradation is a critical design requirement, since RFID infrastructure is subject to real-world disruptions — reader failures, network issues, environmental interference. Rather than returning unavailable errors that exclude the store from agent-directed routing entirely, the service must fall back to the most recently confirmed state with a reduced confidence indicator, signal the degraded state in the response, and trigger an operational alert [21].

#### 4.3 Industry and Societal Implications

The transformation of physical stores into agent service nodes redefines the competitive role of the physical store in agent-mediated commerce [22]. Retailers that have invested in the required RFID infrastructure and API layer will have their inventory surfaced as a real-time fulfillment option in agent queries; those that have not will be treated as unavailable regardless of actual inventory quality [23]. This risk is not purely hypothetical: research on channel integration shows that customers already form their overall retail experience judgments around how consistently inventory and fulfillment information is presented across channels, meaning that a store invisible to an agent is effectively invisible to a growing share of the customer experience itself [12]. There are also equity implications: large, well-resourced retailers are better positioned to implement this infrastructure than smaller independent retailers or regional chains, and if the agent commerce ecosystem develops in a way that systematically advantages large retailers, the resulting market concentration carries negative implications for consumer choice and community-level economic vitality [24]. This unevenness echoes a broader pattern documented in recent work cataloguing the operational, channel, and organizational obstacles that retailers of varying scale face when pursuing omnichannel integration, obstacles that agent-readiness investment now compounds rather than resolves for smaller operators [20]. Industry standards bodies should work toward open, interoperable store inventory API standards that reduce implementation cost and make agent readiness accessible to a broader range of retail operators [25].

### 5. Conclusion

The physical retail store stands at an architectural inflection point. The RFID and operational technology infrastructure deployed over the past decade to improve store operational productivity is also the foundation on which agent-ready store infrastructure must be built. This transition requires deliberate investment in four capabilities: item-level real-time RFID accuracy sustained at a high level, location-aware inventory mapping at a fine-grained resolution, fulfillment channel routing intelligence exposed as a real-time queryable service, and store capacity signaling that enables cross-store fulfillment optimization. The Store Inventory Service proposed in this paper provides the architectural framework for exposing these capabilities through a unified API, and the Store Readiness Assessment Framework provides a practical diagnostic model for evaluating current state against the agent-ready target. As agentic commerce scales, the physical stores that have invested in these capabilities will participate fully in agent-directed fulfillment routing, capturing an advantage no digital-only retailer can replicate.

#### References

1. E. Claucherty, D. Cummins, and B. Aliakbarian, "RFID Unpacked: A Case Study in Employing RFID Tags from Item to Pallet Level," *Electronics*, vol. 14, no. 2, art. 278, 2025. <https://www.mdpi.com/2079-9292/14/2/278>

2. B. Ovezmyradov and H. Kurata, "Omnichannel fulfillment and item-level RFID tracking in fashion retailing," *Computers & Industrial Engineering*, vol. 168, art. 108108, 2022. <https://www.sciencedirect.com/science/article/abs/pii/S0360835222001784>
3. W. C. Tan and M. S. Sidhu, "Review of RFID and IoT integration in supply chain management," *Operations Research Perspectives*, vol. 9, art. 100229, 2022. <https://www.sciencedirect.com/science/article/pii/S2214716022000070>
4. G. Casella, S. Filippelli, B. Bigliardi, and E. Bottani, "Radio frequency identification technology in logistics: A review of the literature," *International Journal of RF Technologies: Research and Applications*, vol. 12, no. 2, pp. 69–86, 2022. <https://journals.sagepub.com/doi/10.3233/RFT-220321>
5. E. Bottani, S. Correale, A. Rizzi, R. Stefanini, and L. Vezzani, "How RFID technology enables the supply chain and retail 4.0: A survey with Italian stakeholders," *International Journal of RF Technologies: Research and Applications*, vol. 14, no. 2, pp. 125–150, 2024. [https://www.researchgate.net/publication/376972419\\_How\\_RFID\\_technology\\_enables\\_the\\_supply\\_chain\\_and\\_retail\\_4\\_0\\_A\\_survey\\_with\\_Italian\\_stakeholders](https://www.researchgate.net/publication/376972419_How_RFID_technology_enables_the_supply_chain_and_retail_4_0_A_survey_with_Italian_stakeholders)
6. H. Siddiqui, F. Khendek, and M. Toeroe, "Microservices based architectures for IoT systems - State-of-the-art review," *Internet of Things*, vol. 23, art. 100854, 2023. <https://www.sciencedirect.com/science/article/abs/pii/S2542660523001774>
7. V. Varriale, A. Cammarano, F. Michelino, and M. Caputo, "Integrating blockchain, RFID and IoT within a cheese supply chain: A cost analysis," *Journal of Industrial Information Integration*, vol. 34, art. 100486, 2023. <https://www.sciencedirect.com/science/article/pii/S2452414X23000596>
8. M. I. Zulfiqar, A. Khalid, A. Siddig, M. J. Nawaz, and S. Saay, "AI-Driven Smart Shopping Carts With Real-Time Tracking and Inventory Forecasting for Enhanced Retail Efficiency," *IEEE Access*, vol. 13, pp. 55576–55585, 2025. <https://ieeexplore.ieee.org/document/10937485>
9. S. Das, R. Ravi, and S. Sridhar, "Order Fulfillment Under Pick Failure in Omnichannel Ship-From-Store Programs," *Manufacturing & Service Operations Management*, vol. 25, no. 2, pp. 508–523, 2023. <https://www.contrib.andrew.cmu.edu/~ravi/msom22.pdf>
10. L. Yang, X. Li, and N. Zhong, "Omnichannel retail operations with mixed fulfillment strategies," *International Journal of Production Economics*, vol. 254, art. 108608, 2022. <https://www.sciencedirect.com/science/article/abs/pii/S0925527322001918>
11. Z. Li, S. Li, and W. Mei, "Buy online and pickup in-store: Co-opetition strategy of omnichannel supply chain players," *Electronic Commerce Research*, vol. 25, no. 1, pp. 369–417, 2025. <https://dl.acm.org/doi/abs/10.1007/s10660-023-09693-6>
12. W. Gao, H. Fan, W. Li, and H. Wang, "Crafting the customer experience in omnichannel contexts: The role of channel integration," *Journal of Business Research*, vol. 126, pp. 12–22, 2021. [https://www.researchgate.net/publication/348170080\\_Crafting\\_the\\_customer\\_experience\\_in\\_omnichannel\\_contexts\\_The\\_role\\_of\\_channel\\_integration](https://www.researchgate.net/publication/348170080_Crafting_the_customer_experience_in_omnichannel_contexts_The_role_of_channel_integration)
13. M. R. Rashidi and S. Chinara, "A Comprehensive Study and Performance Evaluation of Indoor Positioning Methods in BLE and UWB Technologies," *IEEE Internet of Things Journal*, vol. 12, no. 22, pp. 46475–46499, 2025. <https://ieeexplore.ieee.org/abstract/document/11172317>
14. S. Sharma, N. Islam, G. Singh, and A. Dhir, "Why Do Retail Customers Adopt Artificial Intelligence (AI) Based Autonomous Decision-Making Systems?" *IEEE Transactions on Engineering Management*, vol. 71, pp. 1846–1861, 2024. <https://ieeexplore.ieee.org/document/9759231>
15. L. Zhen and H. Li, "A literature review of smart warehouse operations management," *Frontiers of Engineering Management*, vol. 9, no. 1, pp. 31–55, 2022. [https://www.researchgate.net/publication/357848906\\_A\\_literature\\_review\\_of\\_smart\\_warehouse\\_operations\\_management](https://www.researchgate.net/publication/357848906_A_literature_review_of_smart_warehouse_operations_management)
16. J. Kembro and A. Norrman, "The transformation from manual to smart warehousing: an exploratory study with Swedish retailers," *International Journal of Logistics Management*, vol. 33, no. 5, pp. 107–135, 2022. [https://www.researchgate.net/publication/361472763\\_The\\_transformation\\_from\\_manual\\_to\\_smart\\_warehousing\\_an\\_exploratory\\_study\\_with\\_Swedish\\_retailers](https://www.researchgate.net/publication/361472763_The_transformation_from_manual_to_smart_warehousing_an_exploratory_study_with_Swedish_retailers)
17. J. H. Kembro and A. Norrman, "A strategic perspective on automated warehouse systems in retail: insights from a multiple case study," *International Journal of Physical Distribution & Logistics Management*, vol. 55, no. 11, pp. 57–91, 2025. <https://www.sciencedirect.com/org/science/article/pii/S0960003524000059>
18. N. DeHoratius, A. Holzapfel, H. Kuhn, A. J. Mersereau, and M. Sternbeck, "Evaluating Count Prioritization Procedures for Improving Inventory Accuracy in Retail Stores," *Manufacturing & Service Operations Management*, vol. 25, no. 1, pp. 288–306, 2023. <https://edoc.ku.de/id/eprint/27376/1/msom.2022.1119.pdf>
19. Y. Maïzi and Y. Bendavid, "Building a digital twin for IoT smart stores: a case in retail and apparel industry," *International Journal of Simulation and Process Modelling*, vol. 16, no. 2, p. 147, 2021. [https://www.researchgate.net/publication/352776256\\_Building\\_a\\_digital\\_twin\\_for\\_IoT\\_smart\\_stores\\_a\\_case\\_in\\_retail\\_and\\_apparel\\_industry](https://www.researchgate.net/publication/352776256_Building_a_digital_twin_for_IoT_smart_stores_a_case_in_retail_and_apparel_industry)
20. J. Radomska, A. Kawa, M. Hajdas, P. Klimas, and S. C. Silva, "Unveiling retail omnichannel challenges: developing an omnichannel obstacles scale," *International Journal of Retail & Distribution Management*, vol. 53, no. 13, pp. 1–20, 2024. <https://www.emerald.com/ijrdm/article/53/13/1/1267276/Unveiling-retail-omnichannel-challenges-developing>

21. K. Rekulapalli, H. Nilagiri, and M. Vakkanthula, "A Comparative Evaluation of National FHIR Implementation Guides: Strategic Methodologies for Enhancing Interoperability, Regulatory Alignment, and Clinical Comprehensiveness in Global Health IT," *Health Systems*, pp. 1–13, 2026. <https://doi.org/10.1080/20476965.2026.2728621>
22. H. Gaggar, "MemGov: Policy-Governed Memory Lifecycle Management for Enterprise LLM Agents," *Journal of Intelligent Decision Making and Information Science*, vol. 3, no. 11s, pp. 2199–2214, 2026.
23. A. Bhatnagar and V. R. K. Bonagiri, "Operational Excellence through Cloud-Native Distributed Systems," *International Journal of Computational and Experimental Science and Engineering*, vol. 12, no. 2, 2026. <https://doi.org/10.22399/ijcesen.5207>
24. H. Gaggar, "Architecting Scalable Enterprise AI Agent Infrastructure for Distributed Intelligent Systems," *International Journal of Advances in Signal and Image Sciences*, pp. 193–208, 2025. <https://doi.org/10.29284/da9w0117>
25. C. T. Bobba, A. G. Vijay, V. Aher, J. Singh, A. Kejriwal, K. Srinivasan, and A. Ranjan, "Radiographic Evaluation of Periapical Healing Following Single vs Multiple Visit Root Canal Treatment: A PAI Based Study," *Journal of Pharmacy & Bioallied Sciences*, vol. 18, Suppl. 1, p. S206, 2026.