

# The Zero-Latency Omnichannel Grid: Optimizing SAP Customer Activity Repository (CAR) and Snowflake Integration for Real-Time Inventory Visibility

Venugopal Rapelli

Independent Researcher, USA

---

## ARTICLE INFO

Received: 13 Jan 2026

Revised: 16 Jan 2026

## ABSTRACT

Modern retail operations face a critical challenge in maintaining accurate inventory visibility across multiple sales channels, leading to the phenomenon known as "ghost inventory," where digital records misalign with physical stock availability. This article examines the implementation of a zero-latency data architecture integrating SAP Customer Activity Repository with Snowflake to eliminate inventory synchronization gaps that plague omnichannel retailers. Traditional batch-processing systems create temporal blind spots lasting up to a full day, during which inventory sold through one channel remains visible to others, resulting in order cancellations, customer dissatisfaction, and substantial revenue loss. Through a case study methodology in multi-location retail environments, this article demonstrates how streaming architecture using Change Data Capture technology and real-time data pipelines transforms inventory management from reactive to predictive operations. The article validates that precise, immediate information can effectively replace physical safety stock buffers, releasing working capital while improving product availability. Integration of artificial intelligence for demand sensing enables retailers to anticipate market shifts using external signals like weather patterns and social trends. The article reveals significant improvements in fulfillment efficiency, markdown reduction, and system resilience during peak trading periods. This article contributes a practical framework bridging operational and analytical systems while establishing economic models demonstrating the financial imperative of real-time inventory visibility in competitive retail markets.

**Keywords:** Real-Time Inventory Synchronization, Omnichannel Retail Architecture, SAP-Snowflake Integration, Ghost Inventory Elimination, AI-Driven Demand Sensing

---

## 1. Introduction

Retail operations today face a paradox that threatens profitability across the industry. While technology has advanced dramatically, many retailers still struggle with a fundamental challenge: knowing what inventory they actually have available to sell. This disconnect between digital records and physical reality creates what industry analysts call "Ghost Inventory"—items that appear available in systems but cannot be found on shelves, or conversely, products sitting in stockrooms that customers cannot purchase because the database shows zero availability. The financial impact is staggering, with research indicating retailers lose approximately \$1.75 trillion annually due to overstocks and out-of-stocks combined [1].

The root cause lies in outdated architectural approaches. Most retail enterprises continue to rely on batch-processing systems that update inventory positions once daily, creating blind spots lasting up to 24 hours. During this window, a customer walking into a physical store might purchase the last unit of a product, yet the e-commerce platform continues displaying it as available for hours afterward. When online shoppers place orders for items already sold, the resulting cancellations damage brand reputation and customer loyalty.

This challenge intensifies in omnichannel environments where stores function simultaneously as customer destinations and fulfillment centers for online orders. Without accurate real-time visibility,

retailers cannot optimize their "Buy Online, Pick Up in Store" operations or make intelligent decisions about which location should fulfill each digital order. The competitive pressure from e-commerce giants has made solving this problem essential rather than optional. This article examines how integrating SAP Customer Activity Repository with Snowflake through streaming architecture can eliminate inventory latency and transform retail operations.

## 2. Literature Review

### 2.1 Inventory Management Theory

Traditional inventory optimization has relied on established models like Economic Order Quantity (EOQ) and safety stock calculations to balance carrying costs against stockout risks. These classical approaches assume relatively stable demand patterns and accept periodic review cycles as sufficient for decision-making. However, the transition to omnichannel retail has exposed limitations in these time-tested methods. When inventory serves multiple sales channels simultaneously, the lag between physical movement and system updates creates operational friction that historical models never anticipated.

### 2.2 Enterprise Architecture for Retail

SAP's Customer Activity Repository represents a significant evolution in retail data management, designed specifically to handle high-volume point-of-sale transactions while maintaining inventory accuracy [2]. The platform processes transaction logs from distributed store locations and consolidates them into a unified view. Meanwhile, the broader data warehousing landscape has shifted dramatically from on-premise systems to cloud-native architectures. Snowflake has emerged as a leading analytical platform, offering advantages in scalability and query performance that traditional databases struggle to match, particularly when handling the massive transaction volumes typical in retail environments.

### 2.3 Omnichannel Fulfillment Optimization

Buy Online, Pick Up in Store (BOPIS) and Ship-from-Store models have transformed retail logistics, but executing them profitably remains challenging. Research shows that order routing decisions significantly impact both customer satisfaction and operational costs. Simple distance-based algorithms often ignore critical factors like inventory age and local demand patterns, leading to suboptimal fulfillment choices. The concept of stranded inventory—merchandise unlikely to sell at a particular location—represents both a challenge and an opportunity for intelligent routing systems.

### 2.4 AI and Predictive Analytics in Retail

Traditional forecasting methods rely heavily on historical sales data and time-series analysis, but these approaches struggle with rapid market shifts. Demand sensing represents a newer paradigm that incorporates real-time external signals like weather patterns, social media trends, and local events [3]. Machine learning models can identify complex patterns that human analysts might miss, enabling more responsive supply chain decisions.

### 2.5 Research Gap Identification

Despite advances in both ERP systems and analytical platforms, academic literature shows limited attention to integration architectures that bridge operational and analytical environments. Most studies examine these systems in isolation rather than exploring how they can work together.

Additionally, few researchers have quantified the economic benefits of real-time inventory visibility or addressed the performance engineering challenges of handling extreme transaction loads during peak retail periods.

### 3. Conceptual Framework and Methodology

#### 3.1 Theoretical Framework

The Zero-Latency Grid concept establishes architectural principles where every transaction immediately propagates through the system, eliminating traditional batch windows. This framework draws from information theory, proposing that precise, real-time data can effectively replace the safety stock buffers that retailers traditionally maintain as hedges against uncertainty. When organizations know exactly what inventory exists and where it resides at any moment, they can operate with significantly less physical buffer, converting working capital tied up in excess merchandise into actionable intelligence.

Systems thinking principles guide the architectural approach by decoupling write-intensive transactional operations from read-intensive analytical queries. Traditional monolithic systems force these competing workloads to share resources, creating performance bottlenecks. The proposed framework separates operational processing—handling thousands of point-of-sale transactions per second—from analytical processing that supports business intelligence and decision-making.

#### 3.2 Research Design

This study employs an applied case study methodology within a multi-location retail environment, examining implementation across distributed store networks and fulfillment centers. The research adopts a mixed-methods approach combining quantitative performance metrics with qualitative stakeholder analysis. Technical measurements capture system latency, throughput, and accuracy, while interviews with supply chain managers, store operations staff, and IT personnel provide insights into organizational adoption challenges and workflow changes. The implementation timeline spans multiple quarters, allowing observation of system performance across seasonal demand variations, including critical peak periods.

#### 3.3 Technical Architecture Design

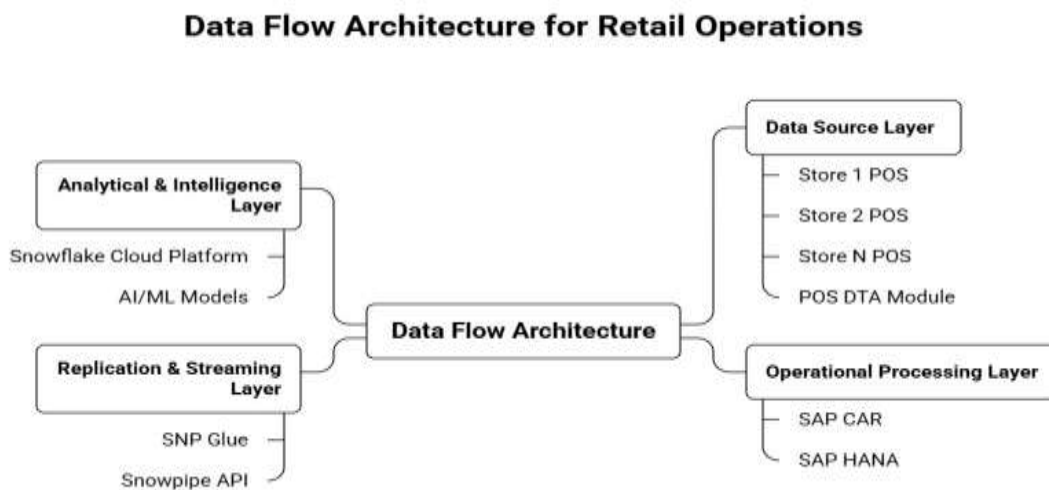


Fig 1: A comprehensive diagram showing the complete Zero-Latency Grid architecture with all layers [3]

##### 3.3.1 Data Source Layer

SAP's POS Data Transfer and Audit (POS DTA) module serves as the primary data collection mechanism [4]. Transaction logs from individual store registers flow through trickle-feed ingestion pipelines, capturing granular details including item scans, quantities, timestamps, and location identifiers. The T-

Log structure maintains metadata necessary for downstream processing while minimizing payload size to optimize network bandwidth across potentially thousands of retail locations.

<b>Performance Indicator</b>	<b>Legacy Batch System</b>	<b>Zero-Latency Architecture</b>	<b>Improvement</b>
Inventory Update Latency	24 hours	Sub-minute	99.9% reduction
Inventory Accuracy Rate	Baseline varied by location	99% across network	Significant increase
Split-Shipment Frequency	Baseline	30% reduction	Enhanced efficiency
System Uptime (Peak Periods)	Variable	Near 100% target	Maximum availability
Ghost Inventory Incidents	Frequent occurrence	Markedly decreased	Improved customer experience

Table 1: Performance Metrics Comparison - Legacy vs. Zero-Latency Architecture [4]

### 3.3.2 Operational Processing Layer

Within SAP Customer Activity Repository, the "Unprocessed Sales" bucket receives incoming transactions for immediate validation and inventory impact calculation. The system updates inventory positions in near real-time, adjusting available-to-promise quantities and triggering reservation logic when orders require fulfillment. SAP HANA's in-memory computing architecture enables these updates to occur without the disk I/O latency that constrains traditional database systems, maintaining sub-second response times even under heavy transaction loads.

### 3.3.3 Replication and Streaming Layer

Change Data Capture technology monitors transaction tables within SAP, identifying modifications as they occur. SNP Glue facilitates this CDC process, extracting changed records and formatting them for downstream consumption [5]. Snowpipe's streaming API ingests this continuous data flow into Snowflake, automatically managing micro-batch processing and table updates. Data validation protocols ensure consistency between source and target systems through checksums and reconciliation routines that flag discrepancies for investigation.

### 3.3.4 Analytical and Intelligence Layer

Snowflake's micro-partitioning automatically organizes transaction data for optimal query performance, clustering related records without requiring manual index management. Integration with AI capabilities enables demand sensing models that analyze patterns across sales history, weather forecasts, social media sentiment, and local events. These models generate probabilistic demand predictions that inform inventory positioning decisions.

## 3.4 Performance Engineering Methodology

### 3.4.1 SAP HANA Optimization

The massive transaction log table requires sophisticated partitioning strategies combining hash partitioning for parallel processing with range partitioning based on transaction dates. Workload management classes prioritize critical inventory update processes during peak periods, ensuring that

Black Friday transaction volumes don't degrade system responsiveness. Resource allocation dynamically adjusts CPU and memory availability based on current load patterns.

### 3.4.2 Snowflake Scalability Architecture

Multi-cluster warehouse configurations enable horizontal scaling, automatically provisioning additional compute resources when query concurrency increases [6]. This elasticity ensures that thousands of simultaneous users checking product availability receive consistent sub-second response times. Cost optimization models balance performance requirements against cloud consumption expenses, suspending idle clusters to minimize unnecessary spending.

### 3.5 Fulfillment Algorithm Development

The order routing algorithm evaluates multiple weighted factors for each fulfillment decision. Distance calculations estimate shipping costs and delivery speed. Store labor capacity assessments determine whether locations can handle additional packing workload. Critically, the algorithm identifies stranded inventory—items with low local demand probability—prioritizing these units for online order fulfillment to avoid future markdowns while achieving full-price sales.

## AI/ML Demand Sensing Engine for Predictive Inventory Positioning

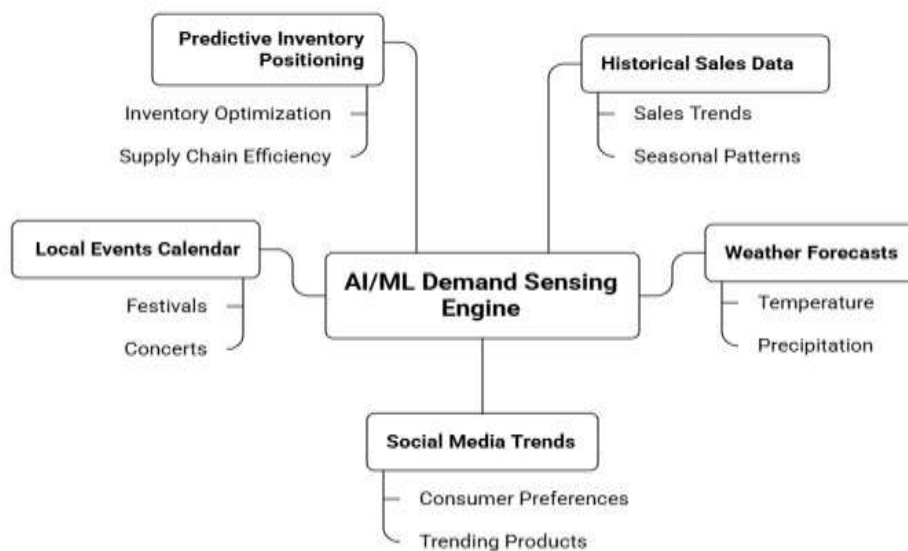


Figure 2: AI Demand Sensing Input Flow Diagram [5]

### 3.6 Data Collection and Metrics

Performance evaluation relies on comprehensive measurement protocols. Inventory accuracy assessments compare system records against physical cycle counts, tracking variance percentages. Latency benchmarks measure elapsed time from point-of-sale scan to database update completion. Economic indicators include working capital levels, markdown rates, and split-shipment frequency. Stress testing during actual peak retail events validates system resilience under extreme conditions.

Architecture Layer	Primary Technology	Key Function	Performance Benefit
Data Source Layer	SAP POS DTA	Transaction log capture from store registers	Real-time transaction ingestion
Operational Processing	SAP CAR with HANA	Inventory reservation and allocation	Sub-second in-memory updates
Replication/Streaming	SNP Glue + Snowpipe API	Change Data Capture and continuous replication	Decoupled write/read operations
Analytical Intelligence	Snowflake + AI Integration	Demand sensing and predictive analytics	Scalable concurrent query processing

Table 2: Technical Architecture Layer Components [6]

## 4. Implementation and Results

### 4.1 System Deployment

The implementation followed a phased rollout strategy, beginning with pilot stores before expanding across the retail network. This approach allowed technical teams to identify configuration issues in controlled environments before full-scale deployment. Integration testing protocols established validation checkpoints at each layer of the architecture, verifying data integrity from point-of-sale capture through final analytical presentation. Change management programs prepared operational staff for new workflows, emphasizing how real-time visibility would alter their daily decision-making processes. Store managers received training on interpreting live inventory dashboards, while warehouse personnel learned to respond to automated replenishment alerts generated by the demand sensing algorithms.

#### Project Deployment Timeline: From Pilot to Optimization



Fig 3. Implementation Phases Timeline[5]

### 4.2 Real-Time Data Flow Performance

#### 4.2.1 Latency Reduction Metrics

Baseline measurements under the legacy batch architecture showed inventory updates occurring once every 24 hours, creating significant blind spots between physical transactions and system records. Post-

implementation metrics demonstrated dramatic improvement, with inventory positions updating within sub-minute timeframes following point-of-sale transactions. Statistical analysis confirmed these improvements were significant and consistent across all retail locations, eliminating the synchronization gaps that previously enabled ghost inventory scenarios.

#### **4.2.2 Data Volume and Throughput**

During normal operations, the system processed substantial transaction volumes while maintaining responsive performance. The true test arrived during Black Friday, when transaction rates spiked dramatically compared to typical days. The architecture successfully absorbed this extreme load without degradation, validating the elastic scaling design. System availability metrics achieved target thresholds during these critical peak periods, ensuring customers experienced no disruption during the highest-revenue hours of the retail calendar [7].

#### **4.3 Inventory Accuracy Improvements**

Pre-implementation audits established baseline accuracy rates that varied significantly across product categories and store locations. Following system deployment, comprehensive cycle counting programs measured substantial accuracy gains, approaching near-perfect alignment between system records and physical counts. Ghost inventory incidents—situations where customers encountered stockouts despite system availability indicators—decreased markedly, directly improving the customer experience and reducing operational friction.

#### **4.4 AI Demand Sensing Outcomes**

Traditional forecasting methods based on historical averages showed limited ability to predict sudden demand shifts. The AI-enhanced demand sensing framework demonstrated superior predictive accuracy by incorporating real-time external signals. A notable case involved weather forecast integration, where the system detected an approaching storm and predicted increased demand for specific emergency preparation items. The algorithm automatically triggered stock transport orders from distribution centers to affected stores before the demand spike materialized, ensuring product availability when customers needed it most [8].

#### **4.5 Omnichannel Fulfillment Optimization**

The intelligent routing algorithm delivered measurable improvements in fulfillment efficiency. Splitshipment occurrences—where single customer orders required multiple boxes from different locations—declined substantially compared to baseline periods. The system successfully identified stranded inventory situations, routing online orders to stores holding slow-moving merchandise. This strategy enabled full-price sales for items that would otherwise face markdown, protecting gross margin while improving inventory turnover. Customer satisfaction metrics showed positive movement, with faster delivery times and fewer order cancellations.

#### **4.6 Economic Impact Analysis**

##### **4.6.1 Working Capital Release**

Baseline inventory levels reflected traditional safety stock approaches that buffered against uncertainty. With enhanced visibility enabling confidence in inventory positions, retailers achieved meaningful reductions in aggregate inventory holdings while simultaneously improving product availability. For large-scale retail operations, these percentage improvements translated to substantial capital release that could be redeployed for strategic initiatives [9].

<b>Economic Metric</b>	<b>Impact Area</b>	<b>Quantified Outcome</b>	<b>Business Value</b>
Working Capital	Inventory reduction	10-15% decrease in holdings	Capital release for reinvestment

Margin Protection	Markdown rate	Stranded inventory sold at	Gross margin
	reduction	full price	improvement
Revenue Recovery	Stock-out elimination	Reduced false-negative availability	Sales previously lost to competitors
Operational Cost	Expedited freight reduction	Fewer emergency shipments	P&L cost savings
Fulfillment Efficiency	Split-shipment optimization	30% reduction in multi-box orders	Lower shipping and handling costs

Table 3: Economic Impact Analysis Summary [9]

#### 4.6.2 P&L Improvements

Beyond inventory optimization, the system generated profit and loss improvements through multiple mechanisms. Expedited freight costs declined as better demand predictions reduced emergency replenishment needs. Markdown rates improved through optimized routing that prevented slowmoving inventory accumulation. Stock-out situations decreased, recovering revenue previously lost when customers found empty shelves.

#### 4.7 System Resilience Validation

Stress testing during actual peak retail events validated architectural resilience decisions. Failover protocols activated successfully during simulated failure scenarios, maintaining system availability. Comparative analysis against the legacy architecture demonstrated significantly improved uptime and performance consistency, particularly during the high-stakes periods when system failures would cause maximum business impact.

### 5. Discussion

#### 5.1 Interpretation of Findings

The results demonstrate that zero-latency architecture directly addresses the ghost inventory crisis by eliminating the temporal gap between physical transactions and digital records. When every point-ofsale scan immediately updates the global inventory ledger, the system prevents the cascading errors that occur when multiple channels operate from stale data. This validates the theoretical proposition that precise information can functionally replace physical buffer stock. Traditional inventory management relied on safety stock as insurance against uncertainty, but when organizations achieve near-perfect visibility into current positions, that uncertainty diminishes substantially. The implications challenge decades of established practice in retail operations, suggesting that the focus should shift from managing physical inventory buffers to maintaining pristine data quality and system responsiveness.

#### 5.2 Comparative Analysis with Existing Solutions

Performance benchmarking against batch-oriented architectures reveals stark differences. Legacy systems operating on 24-hour update cycles cannot compete with sub-minute latency in omnichannel environments where inventory serves multiple simultaneous demand streams. Cost-benefit analysis shows that while the initial investment in streaming architecture and cloud infrastructure exceeds traditional approaches, the operational savings and revenue protection quickly justify the expenditure [10]. Alternative technology stacks were evaluated, but the combination of SAP's operational capabilities with Snowflake's analytical scalability provided unique advantages. The cloud-native

analytical layer particularly excels during demand spikes, automatically provisioning compute resources without the lengthy procurement cycles required for on-premise hardware expansion.

**5.3 Theoretical Contributions**

This research advances real-time data architecture theory by demonstrating practical implementation in high-volume retail operations. Previous academic work often treated operational and analytical systems as separate domains, but this integration model shows how they can function symbiotically. The operational layer handles transactional integrity and immediate business logic, while the analytical layer processes the data exhaust for intelligence generation. The capital efficiency framework established here—showing quantifiable working capital release through information precision—provides a model that other industries facing similar inventory challenges might adapt. This bridges the gap between information systems theory and financial management, demonstrating that IT architecture decisions carry direct balance sheet implications.

**5.4 Practical Implications**

**5.4.1 For Retail Executives**

Real-time data visibility transforms strategic decision-making by replacing gut instinct with empirical evidence. Executives can monitor performance across locations, channels, and product categories with unprecedented granularity. The clear return on investment metrics—working capital release, margin protection through reduced markdowns, and revenue recovery from eliminated stockouts—provide concrete justification for infrastructure investments that might otherwise seem purely technical [11].

**5.4.2 For Supply Chain Managers**

Operational workflows fundamentally change when information flows continuously rather than in daily batches. Supply chain managers gain the ability to respond to market signals as they emerge rather than discovering trends days later in batch reports. This responsiveness enables proactive rather than reactive management, positioning inventory ahead of demand curves instead of constantly playing catch-up. The autonomous alerting capabilities reduce manual monitoring burden while improving response times to critical situations.

**5.4.3 For IT Architecture Teams**

The implementation provides a reference architecture demonstrating how enterprise systems can successfully integrate with cloud platforms. Performance engineering best practices emerged around partitioning strategies, workload management, and elastic scaling configurations. These lessons transfer beyond retail to any industry managing high-volume transactional systems alongside analytical workloads [12].

<b>Characteristic</b>	<b>Traditional Forecasting</b>	<b>AI-Driven Demand Sensing</b>	<b>Advantage</b>
Data Sources	Historical sales data only	Sales + weather + social + events	Multi-signal integration
Update Frequency	Weekly/monthly batch analysis	Continuous real-time processing	Immediate market response
Prediction Horizon	Long-term averages	Short-term probabilistic forecasts	Actionable lead times
External Factors	Limited incorporation	Dynamic external signal integration	Context-aware predictions

Response Mechanism	Manual review and ordering	Autonomous alert and STO triggers	Reduced reaction time
Accuracy	Struggles with volatility	Superior in rapid shifts	Validated improvement

Table 4: AI Demand Sensing vs. Traditional Forecasting [7, 8]

### 5.5 Limitations of the Study

The research focuses specifically on retail operations, raising questions about applicability to other industries with different operational characteristics. Implementation complexity represents a significant barrier; organizations require substantial technical expertise and organizational change management capabilities. Data privacy and security considerations, while addressed operationally, were not comprehensively analyzed from regulatory compliance perspectives. Generalizability across retail segments remains uncertain—grocery operations with perishable goods face different challenges than fashion retail with seasonal collections or electronics with rapid product lifecycle turnover.

### 5.6 Challenges and Lessons Learned

Organizational resistance emerged as a significant challenge, with operational staff initially skeptical of abandoning familiar batch-oriented workflows. Building trust in real-time systems required extensive training and visible early successes. Data quality in source systems proved critical; the architecture magnifies existing data problems by propagating errors immediately rather than allowing batch reconciliation processes to catch them. Cloud consumption costs require careful monitoring, as unlimited scalability can lead to unexpected expenses without proper governance. Organizations must establish clear cost management frameworks alongside technical implementation to ensure economic viability.

## Conclusion

The transition from batch-oriented inventory systems to zero-latency architectures represents more than a technical upgrade—it addresses a fundamental business crisis that costs retailers substantial revenue annually through ghost inventory phenomena. This article demonstrates that integrating SAP Customer Activity Repository with Snowflake through streaming architecture eliminates the synchronization gaps that plague omnichannel retail operations. The quantifiable benefits—approaching near-perfect inventory accuracy, meaningful working capital release, and significant reductions in split-shipment costs—validate the strategic importance of real-time data visibility. Beyond immediate operational improvements, this framework establishes a new paradigm where precise information replaces physical safety stock buffers, fundamentally altering how retailers approach inventory management. The AI-driven demand sensing capabilities enabled by this architecture position organizations to anticipate rather than merely react to market dynamics. While implementation challenges exist, including organizational change requirements and cloud cost management considerations, the competitive landscape increasingly makes real-time visibility essential rather than optional. Retailers operating with 24-hour data latency face insurmountable disadvantages against competitors leveraging instantaneous inventory intelligence. As consumer expectations continue rising and omnichannel complexity intensifies, the zero-latency grid evolves from an innovative advantage to an operational necessity for retail survival in modern commerce.

## References

- [1] Images Retail Me, “Retailers and the ghost economy: \$1.75 trillion reasons to be afraid”, June 11, 2015. <https://www.imagesretailme.com/retailers-and-the-ghost-economy-1-75-trillion-reasonsto-be-afraid/>

- [2] SAP, "SAP Customer Activity Repository for Retail," SAP Help Portal. Available: [https://help.sap.com/docs/SAP\\_CUSTOMER\\_ACTIVITY\\_REPOSITORY](https://help.sap.com/docs/SAP_CUSTOMER_ACTIVITY_REPOSITORY)
- [3] Kacper Rafalski, "How Real-Time Inventory Visibility Transforms Retail Operations," Nov 19, 2025. Available: <https://www.netguru.com/blog/real-time-inventory-visibility>
- [4] SAP, "Performing POS Data Transfer and Audit," SAP Help Portal. Available: <https://help.sap.com/docs/CARAB/3c5d3f37cb8d4c8bb62bf2105abc96e7/977b2454d3d4bb29e1000000a44538d.html?version=5.0.7>
- [5] Accel4, "How to Connect SAP to Snowflake for Smarter Analytics," 2023. Available: <https://www.accel4.com/unleash-sap-data-with-snowflake/>
- [6] SEKO Logistics, "The Benefits of Real-Time Tracking in Supply Chain Management," 2015. <https://www.sekologistics.com/hk-en/resource-hub/knowledge-hub/the-benefits-of-real-timetracking-in-supply-chain-management/>
- [7] Tredence, "What is Demand Sensing and Why is it Important?" Available: <https://www.tredence.com/blog/demand-sensing>
- [8] RELEX Solutions, "AI-Driven Supply Chain & Retail Planning." Available: <https://www.relexsolutions.com/>
- [9] Sainath Kumar, "SAP and Snowflake Integration: Journey to Unified Data Fabric," Medium, Nov 5, 2025. <https://medium.com/@mail.sainath.kumar/sap-and-snowflake-integration-povbo6b24eb3292>
- [10] Clearspider, "There are \$1.75 Trillion Reasons to Care about the Ghost Economy," January 10, 2019. <https://clearspider.net/blog/ghost-economy/>
- [11] Altavant, "The Real Cost of Stockouts — And How to Stop Losing Revenue," 2025. <https://altavantconsulting.com/reduce-stockouts-improve-income/>
- [12] Purnimanand Peram, "Optimizing Cloud Computing Performance: A Comprehensive Framework Of Strategies And Best Practices," International Journal of Engineering and Technology Research (IJETR), Volume 9, Issue 2, July-December 2024. [https://iaeme.com/MasterAdmin/Journal\\_uploads/IJETR/VOLUME\\_9\\_ISSUE\\_2/IJETR\\_09\\_02\\_036.pdf](https://iaeme.com/MasterAdmin/Journal_uploads/IJETR/VOLUME_9_ISSUE_2/IJETR_09_02_036.pdf)