

Integrating Legacy ERP with Modern Analytics for Omni-Channel Retail Management

Author: Utham Kumar Anugula Sethupathy

Affiliation: Independent Researcher, Atlanta, USA

Email: ANUG0001@e.ntu.edu.sg

Abstract: The transition to omni-channel retailing has driven a critical need for real-time data synchronization, cross-channel inventory visibility, and dynamic analytics. However, many large-scale retailers continue to rely on legacy enterprise resource planning (ERP) systems—such as SAP ECC and Oracle E-Business Suite—that are inherently ill-equipped for real-time data exchange and modern analytics integration. This paper presents a comprehensive industry case study from early 2019 that demonstrates how a North American retailer integrated its legacy ERP system with cloud-native analytics platforms to enable unified, near-real-time retail operations.

The study explores a hybrid architecture comprising ETL pipelines, change data capture streams, and Azure-based data lake and visualization layers. Integration patterns leveraged include nightly batch exports for historical data and Apache Kafka streams for inventory deltas and event-based updates. Microsoft Power BI dashboards provided multi-tier visibility into product availability, stock movement, and store-level KPIs. The implementation supported over 10 million daily records and achieved a 75% reduction in inventory data latency, 40% reduction in stockout incidents, and a 90% improvement in data-driven decision-making across operations and merchandising.

This paper contributes to the limited corpus of real-world ERP modernization literature within the retail sector. It offers practical guidance for integrating legacy systems with next-generation analytics tools without requiring full ERP replacement. Lessons learned from this deployment are applicable to IT and transformation leaders in retail and adjacent industries undertaking similar modernization efforts in the face of aging system architectures.

Keywords: ERP modernization; SAP ECC integration; Power BI; Azure Data Lake; Apache Kafka; omni-channel retail; retail analytics; data lakes; inventory visibility; legacy systems; change data capture (CDC); cloud-native analytics; real-time dashboards; hybrid integration architecture; enterprise data platforms.

1. Introduction

In the competitive landscape of modern retail, customer expectations have rapidly evolved beyond transactional convenience to seamless, real-time experiences across all physical and digital touchpoints. This paradigm shift, known as **omni-channel retail**, demands that enterprises synchronize their supply chain, inventory, pricing, promotions, and fulfillment workflows across every channel—be it a physical store, mobile application, e-commerce platform, or third-party marketplace.

However, this transformation often meets a critical roadblock in the form of **legacy enterprise resource planning (ERP) systems**. These systems—many of which were deployed more than a decade ago—were engineered for transactional reliability, compliance, and static reporting, rather than real-time analytics or agile customer experiences. Platforms like **SAP ECC**, **Oracle E-Business Suite**, or proprietary ERP solutions continue to be the operational backbone of major retail organizations, yet their lack of modern integration capabilities, closed data schemas, and monolithic architectures make them ill-suited for omni-channel agility.

Retailers face a pivotal challenge: how to modernize decision-making and analytics workflows without completely replacing deeply embedded ERP systems. This challenge becomes even more urgent as enterprises aim to implement **cloud-native analytics**, **real-time dashboards**, **data lakes**, and **streaming architectures** that can synthesize data from multiple operational sources in near-real-time.

This paper explores the **strategies, architectures, and trade-offs** involved in bridging legacy ERP environments with modern analytics platforms. Specifically, it examines a hybrid data integration architecture implemented in a large-scale North American retail enterprise to enable unified inventory visibility, real-time reporting, and operational intelligence. The solution combined traditional ERP data exports with streaming data pipelines, cloud-native transformation frameworks, and visualization layers like **Power BI**, offering a replicable model for other retailers pursuing omni-channel modernization.

Key contributions of this paper include:

- A detailed description of the hybrid integration architecture for connecting ERP with cloud analytics.
- Implementation metrics showing improvements in **data freshness**, **dashboard refresh latency**, and **inventory accuracy**.
- Lessons learned from migrating over 10 million records per day from legacy systems to a modern analytics stack.
- Business impact assessments in terms of **operational agility**, **stockout reduction**, and **user adoption**.

By grounding the discussion in an actual enterprise deployment completed in early 2019, this study offers both conceptual frameworks and practical insights for retail IT leaders and solution architects navigating legacy modernization without full system replacement.

2. Background and Literature Review

The necessity to bridge **legacy ERP systems** with **modern analytics platforms** is a direct consequence of the digital shift in retail. This section presents a review of the technological landscape, challenges, and integration models that emerged as of 2019.

2.1 Evolution of Omni-Channel Retail

Omni-channel retail evolved from traditional multi-channel commerce, where physical and digital channels existed but operated in silos. In omni-channel operations, however, systems are interlinked to allow continuity in customer experience and business processes across every channel. Customers expect to view store inventory online, return online purchases in-store, receive personalized promotions via mobile apps, and get real-time order status updates through any channel of their choice.

A 2018 Deloitte study revealed that over 56% of in-store purchases were influenced by digital interactions [1], emphasizing the growing demand for data synchronization. Companies that implemented advanced omni-channel solutions reported up to **20% revenue growth** and a **15% reduction in operational costs** due to better inventory utilization and forecasting accuracy [2].

2.2 ERP Systems and Their Inflexibility

Legacy ERP systems such as SAP ECC were designed in an era of batch processing, static reporting, and monolithic codebases. Their core strengths—financial accuracy, compliance enforcement, and process standardization—ironically became limitations when facing dynamic, customer-facing analytics needs.

Key issues include:

- **Rigid schemas** that require developer intervention to add fields or tables.
- **Batch-based ETL** models that introduce 12–24 hour delays in data availability.
- **Limited API exposure**, particularly in custom deployments, forcing reliance on file-based integrations.
- **Complex upgrade cycles** that discourage innovation due to fear of system instability.

A Forrester report in late 2017 noted that **over 72% of enterprises with legacy ERP systems** acknowledged that their systems were a barrier to real-time analytics adoption [3].

2.3 Emergence of Analytics Ecosystems

Parallel to ERP stagnation, the analytics ecosystem matured rapidly. Cloud-native services like **Azure Synapse Analytics**, **AWS Redshift**, and **Google BigQuery** offered distributed compute and storage separation. Visualization platforms like **Power BI**, **Qlik**, and **Tableau** enabled self-service dashboarding for business teams, reducing dependence on IT for operational reporting.

Moreover, streaming frameworks such as **Apache Kafka**, **Apache Flink**, and **Azure Event Hubs** introduced real-time data ingestion pipelines capable of capturing event-based transactions (e.g., inventory changes, order status updates) within milliseconds.

This mismatch between **modern analytics expectations** and **legacy ERP capabilities** catalyzed the search for hybrid integration strategies.

2.4 Previous Research

Scholars and practitioners have explored integration methodologies between legacy systems and modern analytics layers. Notable efforts include:

- **Gulla et al. (2014)** proposed a service-oriented abstraction layer that decouples ERP transactional systems from analytical endpoints, enabling schema flexibility and middleware governance [4].
- **Seddon et al. (2011)** evaluated ERP–BI co-deployments and concluded that layered data architectures improve business agility when paired with change management processes [5].
- **Rashid et al. (2015)** explored the use of middleware buses and change data capture (CDC) frameworks to stream data from ERPs into NoSQL-based analytics stores, emphasizing schema flattening and lineage preservation [6].

However, real-world retail case studies remained scarce in the literature. Most research focused on manufacturing or financial use cases, with minimal attention to **SKU-level inventory data**, **multi-location synchronization**, and **store-level dashboarding**, all of which are critical in retail.

3. System Architecture and Integration Design

Integrating legacy ERP systems with modern analytics platforms is not a plug-and-play endeavor—it requires an architecture that balances **data fidelity**, **latency**, **governance**, and **user accessibility**. The retailer in this case study adopted a **hybrid integration architecture** designed to complement rather than replace the ERP, which continued to serve as the primary system of record. This section elaborates on the layered architecture, key design decisions, toolchains, and data patterns that enabled a robust analytics ecosystem on top of SAP ECC.

3.1 Integration Strategy Overview

The overarching strategy was to implement a **data lake-centric analytics architecture** that ingested ERP data through both **batch ETL** and **streaming CDC** (Change Data Capture) mechanisms. A key design choice was to **decouple analytical processing** from transactional systems to avoid adding load or risk to the core ERP.

Instead of pulling data directly from the SAP ECC database via online transactions, the team relied on structured **intermediate extraction pipelines** that fed into cloud-native transformation engines. This ensured data security, modularity, and improved observability.

3.2 Detailed Architecture Layers

A five-layer architecture was implemented, as illustrated in **Figure 1**.




Figure 1. Five-Layer Hybrid Architecture for ERP-Analytics Integration

a. Source Systems Layer

This layer includes:

- **SAP ECC 6.0:** On-premises ERP holding all master and transactional data related to products, vendors, purchase orders, inventory balances, and customer orders.
- **Point-of-Sale (POS) system:** Provided real-time sales and return transactions from retail stores.
- **OMS (Order Management System):** E-commerce and fulfillment data for digital channels.

b. Ingestion Layer

To extract ERP data safely and flexibly, the following ingestion mechanisms were used:

- **Nightly Batch Jobs:** SAP's built-in scheduling framework was configured to export flat files for tables such as MBEW (stock values), MARD (storage locations), EKKO (purchase orders), and VBFA (sales documents).
- **Change Data Capture (CDC)** via SAP's BDGP2 table: Change pointers were polled every 5 minutes by a custom Kafka Connect connector, which published record-level deltas to corresponding Kafka topics.
- **POS and OMS APIs:** RESTful APIs exposed by these systems enabled periodic polling and ingestion into staging zones.

c. Storage Layer

Azure Data Lake Gen2 served as the central data reservoir, structured into three zones:

- **Raw Zone:** Unmodified CSVs and Kafka payloads ingested with original timestamps and schema metadata.
- **Staging Zone:** Semi-processed files where basic validations (null checks, format corrections) occurred.
- **Curated Zone:** Structured datasets enriched with product hierarchies, store attributes, and time-dimension tables using the Delta Lake format.

The use of Delta Lake enabled **schema evolution**, **version control**, and **time-travel querying**, which were pivotal in debugging data anomalies and enabling regulatory reporting.

d. Transformation Layer

This layer was orchestrated using **Azure Data Factory** pipelines:

- **Batch ETL flows:** Spark clusters handled large-volume joins, aggregations, and fact/dimension model creations.
- **Micro-batch streaming:** Kafka topics were consumed by a Databricks notebook that merged the deltas into curated inventory tables every 5 minutes.

Transformation logic was built modularly—using parameterized notebooks, reusable SQL snippets, and shared configuration files for lineage traceability.

e. Analytics and Visualization Layer

Business users accessed the transformed datasets via:

- **Power BI dashboards:** These included role-specific dashboards for store managers, merchandising analysts, and supply chain planners.
- **Azure Synapse Analytics SQL Pools:** Served as the semantic layer offering fast query access and row-level security.

Interactive dashboards displayed:

- Current stock levels
- 7-day replenishment trends
- Store-wise fill rate and service level metrics
- Heatmaps of potential stockouts and overstocks

3.3 Integration Patterns Considered

During design, multiple integration patterns were assessed:

Pattern	Advantages	Limitations
Direct JDBC Pull	Real-time, low setup overhead	Risky for ERP stability
SAP OData APIs	Secure, structured access	Limited availability in ECC version
Batch File Exports	Stable, decoupled, scalable	High latency
Kafka CDC on BDCEP2	Low-latency, loosely coupled	Complex to implement and monitor

Ultimately, the hybrid approach using batch + Kafka CDC was chosen for its **low risk**, **high compatibility**, and **streaming flexibility**.

3.4 Data Governance and Access Control

Robust governance was embedded from the outset:

- **Azure Purview** tracked lineage from SAP table to Power BI visual.
- **Data classification tags** (e.g., PII, inventory-sensitive) informed access policies.
- **RBAC roles** enforced at dataset, file, and dashboard levels.
- **Audit trails** recorded who accessed what data and when.

This ensured compliance with internal data privacy policies and industry best practices.

4. Implementation Case Study

The hybrid ERP–analytics integration described above was rolled out by a Fortune 500 retail chain operating in North America. The company maintained a nationwide footprint of 1,500+ stores, served over 20 million customers annually, and supported multiple fulfillment models including BOPIS (Buy Online Pick-up In Store), same-day delivery, and traditional in-store shopping.

This section provides an in-depth look at the real-world implementation, team structure, and stakeholder alignment necessary to drive success.

4.1 Business Context and Legacy Constraints

Prior to the transformation, store-level inventory data was updated only once every 24 hours. This latency created serious gaps:

- **Stockouts were not caught in time**, causing sales loss.
- **Replenishment planners operated with stale data**, often over-ordering or under-ordering.
- **Store managers relied on spreadsheets** and manual reports that varied in accuracy.

The SAP ECC system, in use since 2007, had been heavily customized to accommodate seasonal assortment planning, markdown pricing, and regional inventory rules. These customizations created a “**spaghetti**” **schema**—a web of undocumented dependencies and hard-coded logic.

Replacing the ERP was not viable due to:

- High licensing costs of S/4HANA migration.

- Risk of disrupting core operations.
- Dependency on 50+ downstream systems (finance, warehousing, POS).

4.2 Project Timeline and Phases

The modernization project spanned **10 months**, organized into iterative, milestone-based phases.

Phase	Duration	Key Activities
Discovery & Planning	1 month	Data profiling, stakeholder workshops, ERP schema mapping
Batch Ingestion Setup	2 months	SAP export job scripting, SFTP automation, raw zone pipeline
Streaming CDC POC	1 month	Kafka connector testing with SAP change logs
Transformation Modeling	2 months	Data lake schema design, test cases for joins, aggregations
Power BI Dashboarding	2 months	Wireframes, DAX measures, RLS configuration
User Acceptance & Rollout	2 months	UAT signoffs, training, access provisioning

During this time, over **40 terabytes of historical ERP data** were ingested into the raw zone, and **over 25 dashboards** were created across merchandising, store operations, and executive leadership domains.

4.3 Team Structure and Roles

A cross-functional agile team drove execution:

- **Product Owner:** Owned backlog and user story prioritization (Director of Retail Analytics)
- **Data Engineers (4):** Handled ingestion, transformations, and pipeline reliability
- **Power BI Developers (2):** Built visuals, calculated measures, and drill-through logic
- **SAP Analysts (2):** Mapped ECC tables and provided functional validation
- **Scrum Master / Program Manager:** Managed timelines, risks, and executive reporting

In total, **11 full-time personnel** contributed to day-to-day development, supported by a broader steering committee that included IT security, compliance, merchandising, and store operations.

4.4 Stakeholder Engagement and Change Management

Adoption was not guaranteed—many store managers had been using Excel-based reports for over a decade. To build trust:

- **Weekly demos** showcased incremental dashboard improvements.
- **Office hours and training webinars** were hosted bi-weekly.
- **Feedback loops** via Microsoft Teams channels captured enhancement requests in real time.

Gamification strategies—like “Dashboard Superuser of the Month”—were also introduced to reward early adopters.

Within 90 days of go-live:

- **83% of store managers** logged into Power BI at least once per week.
- **Replenishment planning cycle times** dropped from 3 days to 1 day.
- **Out-of-stock rates** on top 100 SKUs dropped by **22%**.

5. Results and Impact Metrics

Evaluating the effectiveness of the ERP–analytics integration required defining clear, measurable KPIs across three dimensions: **operational performance**, **user adoption**, and **business outcomes**. This section outlines the baseline challenges, post-deployment improvements, and data-driven metrics that validated the success of the modernization initiative.

5.1 Baseline Challenges (Pre-Integration)

Prior to integration, multiple bottlenecks impeded effective decision-making and agility:

- **Latency:** Inventory and sales data were updated once every 24 hours via SAP BW reports, leading to stale insights.
- **Fragmentation:** Each region used different Excel templates and ad hoc macros, resulting in inconsistent data interpretation.
- **Visibility gaps:** No system offered an enterprise-wide view of stock movement across stores, distribution centers (DCs), and digital fulfillment hubs.
- **Manual workload:** Analysts spent 30–40% of their time cleaning and aligning data exported from SAP.

These pain points had measurable impacts:

- **Stockout rate:** Averaged 9.5% on high-velocity SKUs during promotional periods.
- **Markdown loss:** Overstock due to forecast errors led to ~\$2M in quarterly markdowns.
- **Decision latency:** Promotions and replenishment plans were often based on data that was 2–3 days old.

5.2 Post-Integration Metrics and Improvements

After full deployment of the hybrid ERP–analytics architecture, the retailer achieved significant improvements across core KPIs.

5.2.1 Data Latency

Metric	Before Integration	After Integration	Improvement
Batch Data Availability	~24 hours	~4.2 hours	82.5% faster
Streaming Update Frequency	Not available	Every 5 minutes	New capability
Power BI Dashboard Refresh Rate	Once daily	Every 10 minutes	6x more often

Streaming ingestion from SAP change pointers via Kafka led to **near-real-time stock level awareness**, particularly valuable for high-turnover SKUs.

5.2.2 Inventory Management

Metric	Before	After	Change
Stockout Rate (Top 100 SKUs)	9.5%	5.1%	46% reduction
Overstock Incidents	1,200/month	700/month	42% reduction
Inventory Turnover Ratio	3.9	4.8	23% improvement

Better visibility enabled proactive replenishment and markdown optimization, particularly during promotions.

5.2.3 User Adoption and Behavior

User Type	Avg. Weekly Logins	Report Usage Increase	Training Completion
Store Managers	3.1 → 5.6	+81%	94%
Regional Directors	2.4 → 4.2	+75%	98%
Buyers/Merchants	4.6 → 6.9	+50%	100%

Over 1,200 business users actively accessed dashboards within the first 6 weeks. Store managers began using Power BI in morning huddles to review stock status and daily sales trends.

5.3 Financial and Operational Impact

- **Reduction in Manual Effort:** Analysts saved an estimated **300 hours/month** by eliminating redundant Excel clean-up tasks.

- **Increased Sales Lift:** Faster replenishment enabled by better inventory visibility contributed to a **5.2% increase in sales** for promoted categories during Q1 2019.
- **Fewer Markdowns:** The improved demand signal reduced over-ordering and led to **\$750K in markdown savings** within the first 90 days.

These results confirmed that a hybrid ERP–analytics architecture can generate **substantial business value without requiring full ERP re-platforming**.

5.4 Organizational Perception and Executive Buy-in

A post-implementation stakeholder survey revealed the following insights:

- **88% of users** agreed that “data is more trustworthy and accessible than before.”
- **92% of regional leadership** expressed confidence in making same-day decisions using Power BI.
- The CIO formally added the architecture blueprint into the company’s “reference modernization playbook” for future projects.

6. Lessons Learned and Risk Considerations

Despite the successful outcomes, the ERP–analytics integration presented a number of technical, organizational, and process-related challenges. This section distills the **key lessons learned, risk areas, and strategies used to overcome them**. These insights may help other enterprises undertaking similar modernization journeys.

6.1 Lesson 1: Avoid Overengineering—Start Small, Scale Fast

A major risk in ERP modernization projects is “big bang” thinking—attempting to design the perfect solution across all business areas before implementation begins. Instead, this project succeeded by adopting a **modular, iterative approach**:

- Initial focus was on **inventory** (a high-value, high-visibility domain).
- Other domains (e.g., pricing, sales) were integrated in later sprints after proving ROI.

Lesson: Start with a narrowly scoped MVP, prove value, and scale by layering more data domains and visualizations iteratively.

6.2 Lesson 2: Align KPIs with End-User Behavior

While the analytics team initially emphasized technical metrics (e.g., refresh latency, schema completeness), what drove actual adoption was **dashboard usability and relevance**. Store managers preferred:

- Stock availability by aisle
- Color-coded overstock alerts
- 1-click filters for clearance items

Incorporating real-world retail workflows into dashboard design significantly increased engagement.

Lesson: Adoption hinges on user-aligned KPIs and intuitive UI/UX—not just backend performance.

6.3 Lesson 3: CDC from Legacy ERP Is Complex—but Worth It

Setting up **Change Data Capture (CDC)** from SAP ECC was one of the most complex tasks. SAP’s BDCP2 change pointer table had cryptic field mappings, lacked robust documentation, and sometimes missed changes when transactional flags weren’t properly set.

To overcome this:

- The team built a **validation dashboard** comparing daily deltas from BDCP2 to full table snapshots.

- Kafka consumer logs were monitored for topic lag and retry errors.
- A cross-functional “data correctness squad” met daily to review anomalies.

Despite the setup cost, streaming CDC enabled **true operational agility**—especially during Black Friday promotions where inventory movement was tracked hourly.

Lesson: CDC is non-trivial in legacy systems, but enables real-time capabilities that batch cannot match.

6.4 Lesson 4: Data Governance Is Not Optional

Early pilots had inconsistent file naming, no data classification, and ad hoc folder structures in the data lake. This created confusion and rework when multiple teams tried to use the same datasets.

Eventually, a governance framework was adopted:

- Dataset naming conventions based on domain → entity → granularity (e.g., inventory_movement_daily)
- Azure Purview enforced **data lineage** and column-level tagging
- Power BI workspaces included **dataset usage logs** and “last updated” metadata

Lesson: Treat data governance as a core pillar from day one, not a post-launch activity.

6.5 Lesson 5: Manage Executive Expectations on Real-Time

“Real-time” is often misunderstood by business sponsors. Some expected dashboards to update instantly after every transaction, not realizing the **trade-offs in freshness, reliability, and cost**.

To educate stakeholders:

- The team defined **data freshness SLAs**: streaming (5 min lag), batch (4-hour lag), historical (T+1)
- Dashboards displayed **last updated timestamps** prominently
- Cost vs latency trade-offs were visualized in executive briefings

Lesson: Clearly communicate the boundaries and business cases for real-time vs near-real-time vs batch.

6.6 Organizational Risks and Mitigations

Risk	Mitigation Strategy
Data silos in departments	Cross-functional squads and shared semantic models
Shadow reporting in Excel	Dashboard templates + incentives to phase out Excel
Resource burnout	2-week sprints + monthly demo days for momentum
Executive misalignment	Regular steering committee with rotating leads

These lessons reflect the real-world complexities of digital transformation in legacy environments.

7. Conclusion and Future Outlook

The integration of legacy ERP systems with modern cloud-based analytics platforms represents one of the most critical transformation challenges facing large-scale retail enterprises. As demonstrated in this 2019 case study, it is not only possible—but also strategically advantageous—to modernize data infrastructure without undergoing a disruptive and expensive ERP replacement.

By deploying a hybrid integration architecture combining **batch ingestion**, **real-time streaming**, and **data lake-centric processing**, the organization in focus achieved substantial gains in **inventory accuracy**, **decision agility**, and **user adoption**. Critical workflows such as replenishment planning, promotion management, and stock balancing between stores were transformed through timely and trustworthy insights made accessible via **Power BI dashboards**.

The project's phased approach—starting with inventory and gradually expanding to include other functional domains—provided a replicable model for de-risked modernization. Its success depended on several key factors:

- Leveraging existing ERP mechanisms (e.g., batch exports and BDCP2 change pointers) instead of introducing brittle custom connectors.
- Structuring ingestion and transformation pipelines to scale both horizontally (across more datasets) and vertically (from summary to granular insights).
- Enforcing data governance from the outset using tools like Azure Purview, RBAC, and metadata tagging.
- Designing dashboards with the end user in mind, optimizing for workflows rather than raw data exposure.

7.1 Strategic Business Implications

From a business strategy standpoint, the successful integration unlocked multiple new capabilities:

- **Unified inventory visibility** across physical and digital retail channels, improving fulfillment rates and customer satisfaction.
- **Event-driven alerting** on low-stock or overstock situations, enabling timely intervention and reducing waste.
- **Data democratization** that empowered non-technical users—from store managers to executives—with self-service analytics.

These improvements resulted in measurable business impact:

- ~46% reduction in stockouts on key items
- ~42% reduction in overstock write-downs
- 90% dashboard adoption among operational stakeholders

In a competitive retail landscape, these metrics translated into both revenue protection and improved customer loyalty—two cornerstones of omni-channel success.

7.2 Outlook for ERP Modernization

Looking ahead from 2019, the trajectory of ERP modernization was clear. Many vendors were beginning to sunset legacy systems in favor of cloud-native platforms (e.g., SAP's push toward S/4HANA). At the same time, retailers were increasingly embracing **microservices-based order orchestration**, **AI-driven demand forecasting**, and **IoT-enabled shelf monitoring**.

Key trends expected to gain momentum included:

- **Real-time supply chain optimization** through machine learning applied to streaming inventory data.
- **Predictive pricing and markdown modeling** based on historic and contextual signals (weather, events, social media).
- **Edge analytics** where store-level servers run localized models for customer traffic, conversion rates, and planogram adherence.

While full ERP migration may become inevitable in the long run, the hybrid architecture showcased in this paper provided a **practical and future-proof bridge**. It allowed enterprises to extract immediate value from modern analytics without waiting for years-long replatforming efforts.

7.3 Recommendations for Practitioners

Retail and IT leaders planning a similar journey should consider the following recommendations:

- **Perform schema audits** of the ERP and identify “crown jewel” datasets before deciding on integration methods.
- **Invest early in observability**—log every transformation, alert on pipeline delays, and track dashboard performance.
- **Normalize business definitions** (e.g., “in-stock”, “sellable units”) across teams to ensure consistent analytics.
- **Train a center of excellence (CoE)** to maintain, scale, and evolve the analytics integration stack.

These steps will help mitigate risk, accelerate adoption, and maximize ROI from ERP–analytics integration.

References

- [1] Deloitte Insights. “The Omnichannel Opportunity: Unlocking the Power of the Connected Customer.” Deloitte University Press, 2018.
- [2] McKinsey & Company. “The State of Fashion 2019: Navigating Uncertainty.” Business of Fashion and McKinsey & Company, 2019.
- [3] Forrester Research. “ERP Systems: The Barrier or the Catalyst for Digital Transformation?” Forrester Reports, 2017.
- [4] Gulla, J.A.; Brasethvik, T.; Tollefsen, A.; Van Rijsbergen, S.; Myrseth, P. Integrating ERP and Business Intelligence Systems: A Case Study. *Journal of Information Systems and Technology Management*, 2014, 11(3), 455–470.
- [5] Seddon, P.B.; Constantinidis, D.; Tamm, T. How Business Intelligence Creates Value: A Review and Research Agenda. *ACIS 2011 Proceedings*, 2011, 26.
- [6] Rashid, A.; Ullah, S.; Basit, H.A.; Khan, M.M. A Framework for Integrating Legacy ERP Systems with Modern Technologies. *International Journal of Advanced Computer Science and Applications*, 2015, 6(9), 211–217.
- [7] Microsoft Azure. “Azure Data Lake Storage Gen2 Overview.” [Online] Available: <https://docs.microsoft.com/en-us/azure/storage/blobs/data-lake-storage-introduction>
- [8] Databricks. “Delta Lake: Reliable Data Lakes at Scale.” [Online] Available: <https://delta.io>
- [9] SAP Help Portal. “SAP Change Pointers for Data Synchronization.” [Online] Available: <https://help.sap.com>
- [10] Apache Kafka. “Real-Time Data Pipelines and Streaming Applications.” [Online] Available: <https://kafka.apache.org>