

MongoDB for Retail

Essential elements to
ecommerce modernization

Table of Contents

Executive Summary: Ecommerce Modernization and Opportunities to Transform	4
Building a Modern Product Catalog	4
Modernization Drivers for Product Catalog	5
Application-Driven Analytics in Retail	7
Build Deeper Business Insights with In-app Analytics and Real-time Business Visibility	7
In-app Analytics	8
Real-time Business Visibility	8
Personalization and Search	9
Personalization for Online and In-person Shoppers	9
Easy-to-build Retail Search Experiences	10
Worlds Collide: Search Personalization	11
The Single View: Constructing Customer 360	12
The Trick to Building a Single View	12

Table of Contents

Getting Loyalty Programs Right with Receipt Mining	13
Sales Receipts, the Marketing Gold Mine for Retailers	13
Single View: the Trick to Building a Concrete Customer Loyalty Program	14
Achieving Greener Commerce with Real Time Stock and Inventory Management	14
Transforming Product and Inventory Management Through Automation	15
Enabling Mobile Solutions and Omnichannel in the Retail Value Chain	16
Workforce Enablement: Why It's Worth the Investment	16
Challenges to Building Your Retail App	17
Make Your Data Accessible All the Time, Even Offline	18
Legal Notice	18
The Next Steps to Modernization	19
Resources	19



Executive Summary: Ecommerce Modernization and Opportunities to Transform

Global retail ecommerce sales are expected to grow 56% by 2026, reaching \$8.1 trillion, accelerating retailers' appetite to pick up the pace of innovation, focus on customer experience, and secure long-term business prosperity.

With MongoDB, enterprises are enhancing ecommerce experiences and preparing for the growth ahead with better product catalogs, search, and recommendations while slashing costs with real-time analytics in the supply chain, ushering modern retailers into the next generation of online selling.

Learn how MongoDB Atlas is enabling retailers to move to the forefront of ecommerce modernization and delight customers with application-driven insights through a unified, fully-managed, cost-effective developer data platform. Dive into this paper to traverse seven key pillars to providing experiences modern retail consumers have come to love and expect from the businesses they trust.

1. Building a Modern Product Catalog
2. Application-Driven Analytics in Retail
3. Personalization and Search
4. The Single View: Constructing Customer 360
5. Getting Loyalty Programs Right with Receipt Mining
6. Achieving Greener Commerce with Real Time Stock and Inventory Management
7. Enabling Mobile Solutions and Omnichannel in the Retail Value Chain

Each of these chapters covers the essentials you need to get started, and provides additional resources to delve into when you and your team is ready. Let's dive in!

Building a Modern Product Catalog

Product catalog data management is a complex problem for retailers. Older, more established retailers have traditionally relied on vendor-supplied product information management systems (PIMS) as their backend data store.

The PIMS is a tool or set of tools built specifically to help retailers produce and manage product content across the various digital channels they support. Typically, this will act as the "single source of truth" across the organization for product information, including technical data, usage data, emotional data, and digital catalog data (number of products, variations, seasonality).

Product data will be input to the PIMS and then a subset of it will be extracted, transformed, and loaded to a product catalog data store to be combined with other necessary information that

the viewer will require. For example, the product catalog for ecommerce will need review data and upsell recommendations, whereas the product catalog for the workforce mobile app will need in-store stock location information. The rigidity of the relational database management systems (RDBMS) used often meant that a new product catalog data store was created on top of the PIMS for each use case. This led to a reduction in time-to-market for new products, data hygiene and duplication problems, and long delays for product catalog updates.



After years of relying on PIMS and multiple monolithic, vendor-provided systems, retailers have learned that product catalogs built on legacy databases are unsuitable for modern ecommerce experiences. Retailers are now increasingly approaching their product catalog modernization by implementing an operational data layer (ODL) that centrally integrates and

organizes siloed product data, making it available to all consuming applications. This enables a range of use cases such as single view of product, real-time analytics, and omnichannel experiences. Here we'll discuss the modernization drivers, the journey to implement an ODL and why MongoDB's document model and architecture make it a great fit for the use case.

Modernization Drivers for Product Catalog

Broadly speaking, some of the challenges that arise from the PIMS approach that retailers are using today include data silos and the complexity of keeping them consistent and up to date. This causes a lack of access to rich, consistent product

information across sales channels, which is vital for a smooth customer journey, especially in omnichannel experiences that mix mobile, in-store, and online channels.

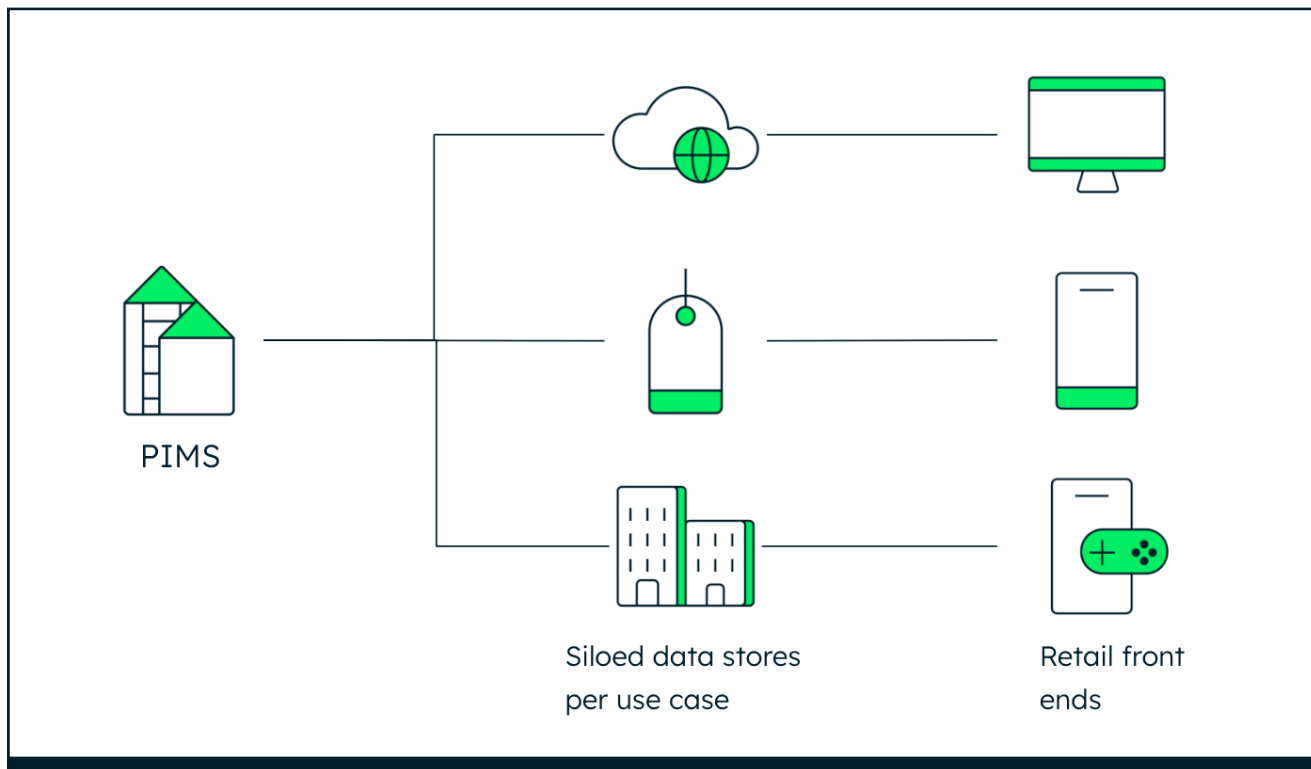


Figure 1. Current PIMS approach shows siloed data stores

Common challenges with legacy product catalogs

- Loss of revenue and customer trust due to inconsistent view of product
- Loss of competitive edge due to delays in new products and info going live
- Increase in churn or bounce rate due to poor product search capabilities
- Inability to increase revenue through cross sell or upsell
- Difficulty maintaining profitability due to low margins and high operational cost

What does a modern product catalog need?

- Single consolidated view of the holistic product catalog
- Flexibility to model and evolve diverse products over time
- Scalability to add products and product data without sacrificing speed
- Search that is easy to enable and provides fast sophisticated results
- Analytical capabilities to offer personalized up sell and cross sell

[See how OTTO manages](#) new catalog updates for two million products.



These are the tenants of a modern operational data layer (ODL), a single, consolidated view across all data that has the flexibility, scalability, and technical capabilities to serve the operational and analytical needs of the consuming business applications.

To learn more, [download our white paper: *Implementing an Operational Data Layer for Product Catalog Modernization*](#) to gain access to our sample reference architecture.



Application-Driven Analytics in Retail

Build Deeper Business Insights with In-app Analytics and Real-time Business Visibility

The basics: application-driven analytics can leverage the same datasets to meet two broad use cases:

- **In-app Analytics:** This enables leveraging real-time data to enhance user experiences and drive immediate user or app actions, resulting in truly intelligent applications providing unique user experiences such as personalized offers and real-time order tracking. Developers need to infuse analytics on live data directly within the operational flow of the application to make this possible.
- **Real-time Business Visibility:** This takes the form of a set of processes leveraging live data to drive actionable insights for the business in real-time. It enables strategic decision-making and process optimization, such as automatically distributing products in a freight container based on inventory and properties of different products to maximize space utilization and avoid damage. Analytics teams are tasked with generating insights without causing any disruption to the application.

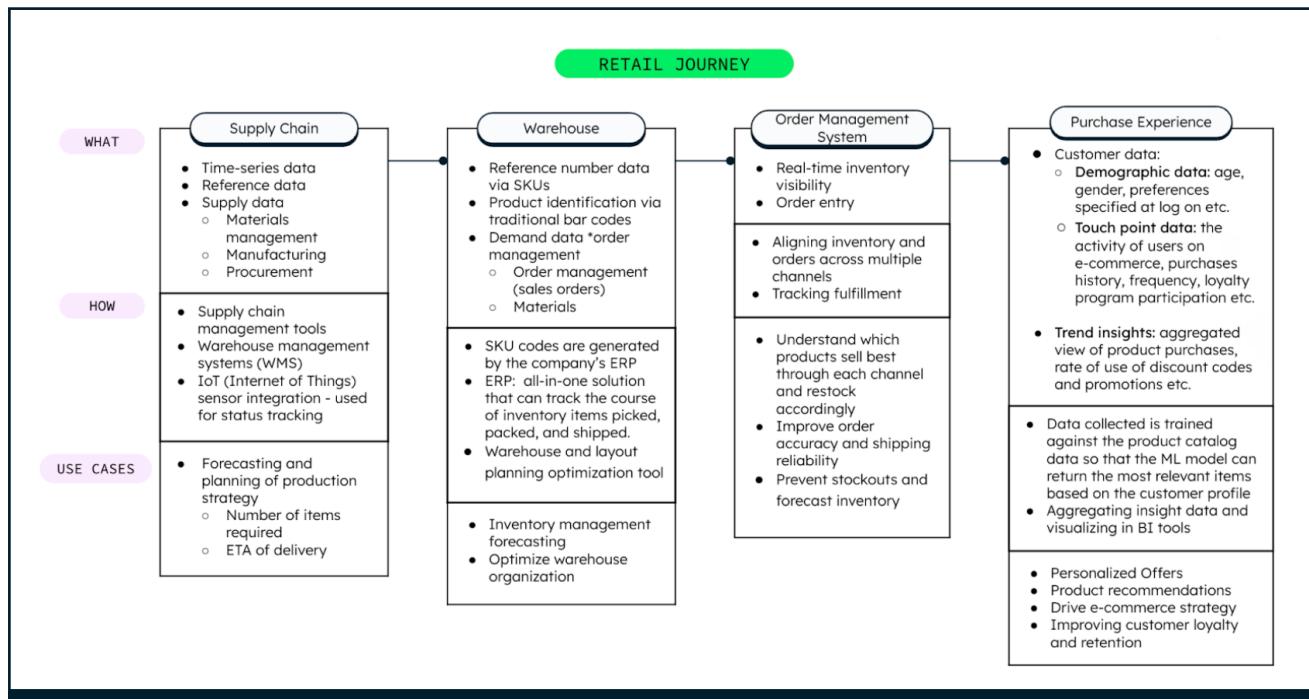


Figure 2. Analytics can augment every step of the retail journey




In-app Analytics

In-app analytics can be leveraged to enable tailored experiences for customers in real-time to provide consistency across physical and digital channels. This allows companies to improve brand visibility, increase conversion rates, and maximize spending. For instance, in-app analytics can be integrated into an e-commerce portal to provide personalized recommendations to customers in real time based on what a customer is doing right now – either on the retailer’s website or in the store. Then the same datasets can be leveraged to adjust pricing in real time, based on individual and environmental factors such as customers’ propensity to spend, how much competitors are charging, and the product’s demand.

Analytics can also play a pivotal role in [enabling your in-store workforce](#) to be more efficient and satisfied. This might translate into an intelligent mobile POS system allowing customer agents to easily recommend in-stock products that

are likely to suit a particular customer based on demographics and previous purchases. You can go a step further and allow them to manage out-of-stock items by scheduling them to be delivered to the customer’s preferred address, with clear ETAs of when they should expect it based on real-time courier data. An intelligent workforce management system can be provided to your staff and HR teams to be able to meet demand surges and needs in different parts of your store, in real time (e.g. at tills, restocking, and warehouse). The opportunities are limitless and can directly impact the employee’s efficiency and retention while helping improve your profit margins.


These are only some of the scenarios where analytics can be leveraged. It can support many other areas such as supply chain, warehouse management, fraud detection, and marketing campaign planning.

[Explore Expedia’s journey](#) to personalization with MongoDB. 

Real-time Business Visibility

Real-time business visibility enables data sitting in operational systems and application-generated data to be combined in real time to make just-in-time strategic decisions. This data can also be leveraged by business intelligence tools (e.g. Tableau, PowerBI) to provide input for centralized analytics systems accessing data from Enterprise data warehouses. The major difference with Enterprise data warehouses is that data generated through applications often take several hops (and time) before making its way into a data warehouse before it is queried.

For instance, building an aggregated view of customers based on key metrics such as demographics, spending habits, and geography can help marketing teams assess the effectiveness of promotion campaigns and analyze their impact on revenues and warehouse in real time (e.g. to terminate a campaign earlier if items become out of stock). Retailers can go a step further and predict the best time to run a flash sale by monitoring trends in the sales pipeline and spotting the main revenue streams for the business (e.g. online vs in-store, sub-brands, age groups).

[Learn how BOXED \(B2B, B2C\) leveraged MongoDB Charts](#) to get consolidated insights for their data warehouse teams to meet 30x demand during COVID-19. 

Dive in deeper to in-app analytics: read the full white paper [here](#).



Personalization and Search

Thanks to early advances by ecommerce behemoths like Amazon, modern consumers expect all of their shopping habits and browsing history to pay into future online experiences.

In the past, retailers knew their customers' names, phone numbers, and addresses. Now ecommerce dealers have access to endless data that not only helps build customer profiles, but tells sellers exactly how to predict consumers' buying habits.

There are two vital and interconnected ecommerce use cases that will make or break the online buying experience for shoppers: personalization and search.

Personalization for Online and In-person Shoppers

Retail ecommerce personalization engines have the ability to build customized online experiences for your customers in real time. Based on analysis of behavioral and demographic profiles, historical interactions, and past preferences, these engines are built on top of legacy customer data management systems. They could also replace legacy systems.

Retailers have faced pressure from discounters and online upstarts, all while reconfiguring their businesses after the pandemic. As consumers return to in-person shopping experiences,

personalization is going omnichannel. Customers are starting to expect personalized store greetings, location-based offers on nearby products, and the hyper-customization of products, services, and special offers.

With in-store beacons, electronic shelf labels, smart mirrors, AI, and a host of other connected technologies, the next frontier of personalization is the seamless connection, and blurring, of the digital and physical retail experience. And customers are expecting this omnichannel, personalized experience to happen in real time, too.

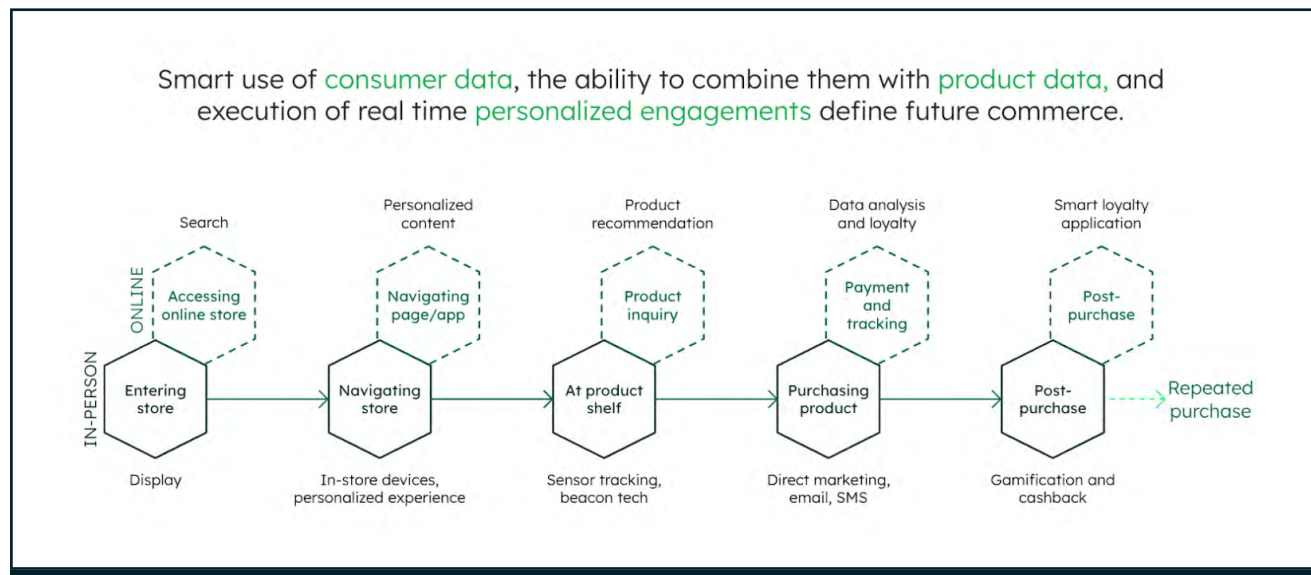


Figure 3: Customers expect personalization. Here's how to use consumer data for personalization across channels.

[Learn more](#) in the MongoDB Guide to Personalized Retail.

Easy-to-build Retail Search Experiences

Just as Amazon has created the standard for ecommerce personalization experiences, Google has transformed customer expectations in search. Highly advanced capabilities like faceting, fuzzy matching, and auto-correct are the basic standards expected by shoppers.

Whether shoppers are browsing for fun or looking to buy a specific item, the search bar is the first place that 87% of consumers begin their product search journey. This means that entire businesses

depend on the reliability of the search bar. In most cases, conversion rates nearly doubled from customers who used on-site search and easily found what they needed. And as for Amazon, its conversion rate shoots up 6x when visitors do a search, from 2% to 12%.


Search is important at various different customer touch points and portions of the supply chain. From the factory floor to a customer's doorsteps, a good search experience drives sales growth.



Workforce apps

Customers and staff alike need the ability to search and quickly find products, store, and customer data.

[7-Eleven](#) uses MongoDB Atlas Search in its internal web portal to guarantee real time inventory data for its brick and mortar stores. Store managers use fuzzy search for internal inventory tracking and to browse internal reports that track sales data.



Inventory and assortment

Decisions about what stock should be in what store, inventory management and categorization of stock all require search for efficiency.

Customers aren't the only people that can benefit from better search tools. MongoDB Atlas Search helps companies build internal web portals where operators can easily find merchandise and products with some typo tolerance, to keep business moving.



Customer 360

An operator on a call line and in-store staff alike need to be able to immediately identify a customer based on a variety of information.

Order numbers, loyalty numbers, phone numbers and email are all important customer data that retailers should be able to quickly search to improve customer experiences. Using MongoDB Atlas Search, retailers can build in-store customer lookups for rewards programs.



Parts and assemblies

Manufacturing parts and assemblies with difficult permutations are challenging to find in large sets with variation.

Automobile makers, dealers, and garages around the world will need to use modern applications to store and browse parts catalogs, repair guidelines, and supply chain data. By using search, dealers can more easily find and leverage upsell and cross-sell opportunities.





Chatbots

It's essential that chatbots can identify customers, orders and products in order and with complete accuracy.

Quick parsing of customer data brings the efficiencies that quickly scaling retail enterprises need to answer queries quickly and precisely.

Effective chatbots answer questions to customers freeing up hours of work time for customer operations teams.

Worlds Collide: Search Personalization

So you've built a great search experience. Don't stop there – modern search experiences will be personalized to your past behavior and additional environmental factors.

In order to get all the right information from customers to enhance search personalization,

retailers must gather data from internal systems, such as the CRM, loyalty programs, stock, and inventory, as well as external data like weather and social media. This works best when data is moving in real time to give the most relevant experience.

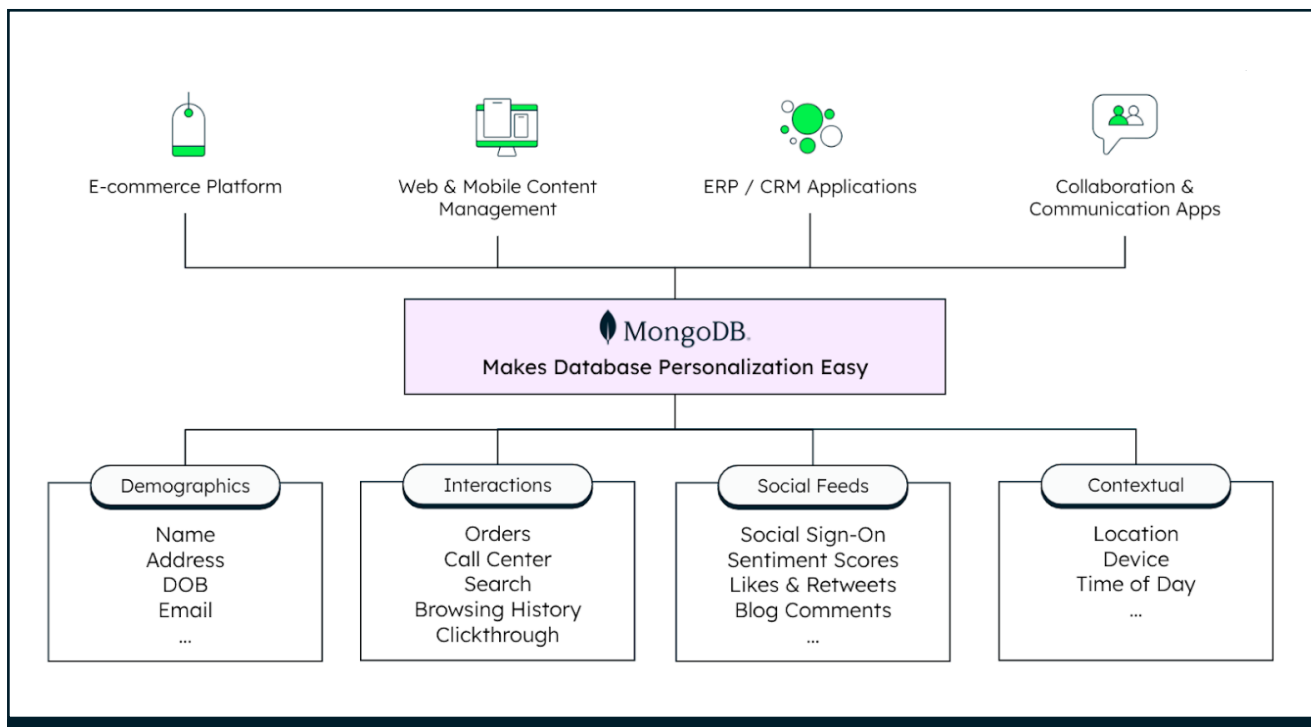


Figure 4: Retailers can greatly benefit from personalized search across various platforms, applications, and data management systems.

Dive deeper with our [white paper: Transforming Customer Experience with MongoDB Atlas Search.](#)



The Single View: Constructing Customer 360

A retail customer 360 establishes a complete view of customers by consolidating data from every customer interaction with a retailer. These can be online, offline, in real time, or next-day for purchase of products, services, or support.

- A customer single view solution *centralizes and organizes* data from all related sources into a standardized format and unified data model
- It addresses the problem of *disjointed and duplicate data* to give a single version of truth of customers
- The single view includes a consumption component providing a *unified view of customers* across all related data domains
- It enables organizations to *rapidly evolve* their schema and aggregation mechanisms
- It serves as an ideal foundation for *analytics, AI/ML, reporting* and enhancing current operational capabilities

The Trick to Building a Single View

It's a big lift for IT teams to bring together different systems storing different types of data from different departments. Most of these systems don't communicate with each other, and bringing them into a single view is very difficult.

Reconciling the schemas of data from different systems is hard, and in some cases impossible. The relational databases often employed by retailers weren't built for diverse data types. In addition, they also struggle to give retailers the agility they need to iterate on their schemas. With rigid legacy schemas, you're pulling in data from legacy customer service systems when tomorrow you may need to incorporate customer sentiment analysis.

Are you giving a feature trade-off? A single view application is only as good as its ability to serve up flawless access to the data within it. You can't sacrifice data access capabilities like ad hoc queries, secondary indices and the ability to collate data for flexibility.

MongoDB makes building a single view of your customers easy with the document model, dynamic schemas, and its expressive query language. With JSON documents, all basic data types are supported without the need to define or enforce a schema. Pull in new data when you need to with MongoDB's self-describing schemas. And with the expressive query language, MongoDB's indexing and aggregation capabilities make it possible to find and filter whatever data the business needs.

So what is the next step to making the most of your customer data? Many retailers are managing customer data and loyalty data that is separate, or not fully integrated. Read on to see how to make the most of this data to build successful loyalty programs.

[See how](#) AO.com built a single view in three months to serve call center, fraud, and GDPR apps.



Getting Loyalty Programs Right with Receipt Mining

Loyalty programs, sponsored by retailers and other businesses, offer rewards, discounts, and other special incentives as a way to attract and retain customers. They are designed to encourage repeat business, offering people a reward for brand loyalty.

A retail customer loyalty platform consists of various aspects of customer interaction at all touch points with a given retailer, where customer shares their preferences, purchase history, opinions and other details in return for accumulating points and receiving discounts that can be used to redeem for various forms of rewards. They are designed to enhance repeat business and encourage brand loyalty.

Customer data associated with loyalty programs can enable numerous analytical capabilities useful for any retailer, including:

- Customer behavior patterns isolating declared and undeclared preferences
- Customer choice of store locations and channels for shopping (online, mobile, in-store, etc.)
- Order patterns with a given retail location/ brand etc.
- The products viewed at various locations: Online, in-store, mobile apps
- What was viewed compared with what is purchased on which channel
- Lifetime spend of a given customer and patterns of transactions per time window (holiday times compared with other times).
- Average spend by channel (who belongs to your loyalty program and who doesn't), by kiosk (identify popular checkout location, etc.)

Sales Receipts, the Marketing Gold Mine for Retailers

The sales receipt is one of the most important data sets in the retail industry. A sales receipt encapsulates a wealth of information on customer shopping habits, product movement, channel information, proof of purchase, promotion effectiveness, as well as loyalty programs customers are engaged in.

While sales receipt data can provide numerous analysis data sets that are extremely useful for any retailer, the disjointed nature of receipt data can make it challenging to mine the most important data for understanding consumer habits.

Sales receipt data contains key information needed by retailers to enable personalization, product recommendation, inventory movement patterns and others when available at the right time.

- Retailers maintain multiple copies of the same sales data which becomes disjointed.
- By nature, ETLs are batch jobs; data is out of date the moment it's available; no real-time capability.
- Machine learning models that use stale data do not reflect current consumer behaviors.

Additionally, systems of record are based on legacy data technologies, which are typically rigid, fragile systems that hold back innovation. There is a lack of access to historical data and related data sources at once to address business needs. These systems lack the capabilities to handle the volume and variety of requests they receive, and are simply unsuitable for modern data-centric application designs.



Single View: the Trick to Building a Concrete Customer Loyalty Program

Creating a customer single view solution will centralize and organize all required data from related sources needed to achieve a modern loyalty program. By moving data into a standardized format and unified data model, the duplicated data sources are moved into a single version of the truth for customers.

The single view, as described in the previous chapter, enables organizations to more quickly evolve their schema and aggregation mechanisms to serve as an ideal foundation for analytics, AI and machine learning, and reporting to enhance current operations.

In 2020, Marks & Spencer (M&S) relaunched its digital-first loyalty scheme, promising its seven million customers a more personalized and tailored experience. With an anticipated growth in adoption, the digital and data team made the bold decision to build the platform in-house (a cultural shift from its traditional 'buy' model). Follow the link below to find out why and how this decision was made, and how the M&S team used technologies like MongoDB to scale the platform to support a 200%+ growth in less than a year.

[VIDEO](#): How Marks & Spencer used MongoDB to scale Sparks, its digital-first loyalty scheme.



Achieving Greener Commerce with Real Time Stock and Inventory Management

Modern customers are not just expecting the highest standards of service and tech in their retail shopping experiences. They also expect a higher standard of sustainability and are aiming to lower their carbon footprint.

There are a plethora of reasons why retail can be a wasteful industry. Between grocers throwing away spoiled food, and the fast fashion trends that clothing dealers have to keep up with to compete, it takes a lot of extra effort for retailers to achieve greener commerce in the supply chain. With the help of better data management and analytics, retailers can streamline the supply chain and find solutions that lead to sustainability.

Despite these obstacles, a great example of a sustainability-driven MongoDB customer comes directly from the retail industry.

Since it began, [Vestiaire Collective](#) has persuaded its consumers to reconsider their typical buying

habits by using resale as a sustainable alternative to fast, throwaway fashion. Through its online ecommerce platform and growing community of 11 million members, Vestiaire Collective has a community that spans more than 80 countries, generating more than 550,000 new listings a month.

As Vestiaire Collective's growth exploded, so did the products, authentications, users, and purchases and additional data sources that the engineering team had to analyze, manage, and store. This led to business challenges like delays to authenticate orders, which put a dent in the customer experience and used up more resources. An upgrade to the business' core technology infrastructure was necessary to improve efficiency.



A switch to MongoDB Atlas allowed the enterprise to manage its growing product catalog, evolving shipping requirements and utilize the ease and scalability of the document model. The result was more than just better customer experiences: Vestiaire Collective also began reducing its

carbon footprint through MongoDB Atlas implementation, which consumes less energy compared with its previous data management system that ran on aging servers. Now the enterprise has reduced its overall server footprint.

Transforming Product and Inventory Management Through Automation

Achieving better sustainability and efficiency in retail means attacking every aspect of the supply chain. From product catalogs to the factory floor, automation helps achieve carbon reduction and sustainability goals.

[Rent The Runway](#), an ecommerce enterprise that allows its customers to rent luxury garments for special occasions, tackled its own sustainability goals while working to streamline complex inventory management challenges. In order to process orders more quickly and efficiently, while also extending garment lifecycles, Rent the Runway invested in robotic arms to process orders and route items through the warehousing and cleaning process.

Rent the Runway has two U.S.-based fulfillment centers that handle the processing of new garments, cleaning, and dispatch to new customers as quickly as possible. The cleaning process is the most complex of these operations. There are dozens of cleaning methods to use depending on an item's color, fabric, and solvent needed. Each garment also requires a check for

forgotten items, like lipstick or ticket stubs left in a pocket. In the past, these checks were performed manually by hand.

Now, the company has streamlined operations by building six MongoDB Atlas clusters for different operations in the fulfillment centers. This allowed the company to move swiftly while also extracting data for data scientists to analyze to fine tune processes.

Today, Rent the Runway saves resources by processing garments through x-rays and sorting them with robotic arms that sort items based on RFID tags. The automation doesn't stop there – after items are organized into bins, the bins send alerts that detect problems like items left in pockets. The data events are all automatic and feed into database logic.

By implementing these changes, Rent the Runway has achieved a 67% decrease in processing times. This gets garments back on the rack faster, giving more customers the opportunity to rent luxury items and save from buying a garment they may realistically only wear once or twice.



Enabling Mobile Solutions and Omnichannel in the Retail Value Chain

Enabling retail workers with real-time data via a mobile app can drastically improve employee productivity, happiness, and retention, and at the same time transform the customer experience employees provide – all factors that impact your bottom line.

Workforce Enablement: Why It's Worth the Investment

If you're not enabling your workforce, you're losing money. Take, for example, employee turnover. The average turnover rate in the retail industry is slightly above 60%, according to the National Retail Federation. This high turnover rate translates into more than 230 million days of lost productivity and \$19 billion in costs associated with recruiting, hiring and training, according to Human Resources Today.

When employees have the digital tools to help them provide a great customer experience and make their own jobs easier, you mitigate problems like burnout and job dissatisfaction. Today's retail workers are digital natives, with access to unprecedented levels of technology. You need modern, digital enablement for the modern worker.

A well architected and offline-first worker app will help your staff, and your business, accomplish key objectives:

- **Improve the customer buying experience:** Frontline staff that are equipped with a mobile-first app are armed with real-time information, like what's actually in stock, and can also make suggestions based on customer buying history.
- **Increase employee productivity:** According to Deloitte, workers spend as much as three hours of their time each week looking for information they need. Imagine the impact regaining those hours could have on worker productivity!
- **Track and improve the performance/sales/buying experience through data analysis:** The potential of workforce enablement apps extends beyond just identifying what items are in stock at which stores. They can also gather valuable data to reveal key patterns in everything from customer purchase habits and peak shopping times to individual worker metrics, like the number of successful sales. With those data insights, you can better allocate workers, assign workers based on strengths, stock items based on buying trends, and more.



Challenges to Building Your Retail App

A modern, innovative retail workforce enablement app sounds great. But for already strained IT teams, building this kind of intuitive app from the ground up is often challenging.

Gaining a single view of your data is essential, but complicated.

Accurate and up-to-date data is the key to everything, from real-time inventory views to personalized customer interactions. It's also the information backbone that feeds powerful mobile workforce apps. But the volume and variety of data across omnichannel shopping experiences is difficult to collect and process in real time, especially when many retailers are using legacy systems that are decades old and make it nearly impossible to have a single source of truth.

A typical retail technology stack, assembled over decades, relies heavily on relational databases, backed by several supporting database, data warehousing, and analytics technologies all trying to ingest and process data from across the business.

The data from all these systems ends up compromised, requiring time consuming ETL maneuvers to bring it together in one place or move it in and out of separate analytics suites. Without a single source of truth, built on a flexible data foundation, meaningful business insights and real-time analytics capabilities are very hard to achieve. Your mobile workforce app displays out-of-date information, or worse, no relevant information at all.

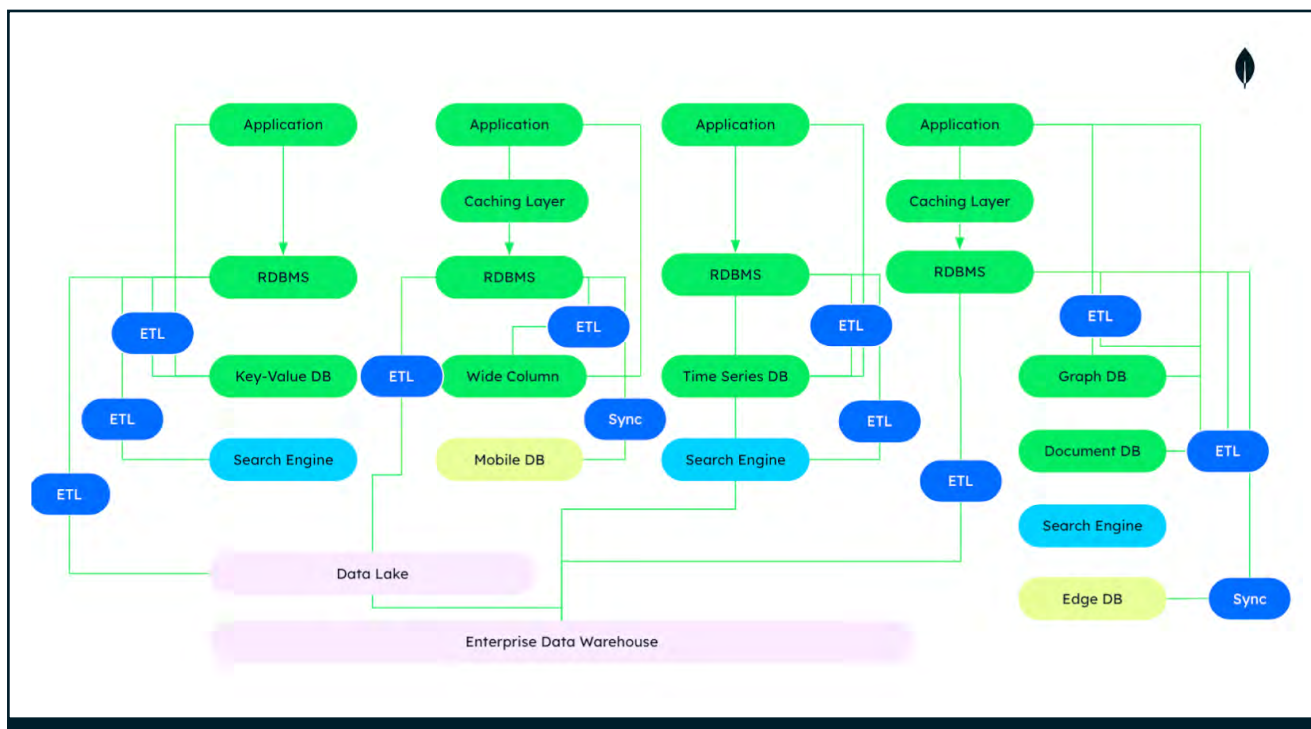


Figure 5: As shown in this typical retail tech stack, legacy architectures are often made up of specialist NoSQL and relational databases, and additional mobile data and analytics platforms – all resulting in siloed data, slow data processing, and unnecessary complexity.



Make Your Data Accessible All the Time, Even Offline

In the fast-paced rhythm of grocery retail operations, shelves are restocked continuously and the prevalence of online ordering and in-store shopping assistance makes the availability of real-time data accessibility essential.

For one MongoDB customer, real-time data service was the key to providing real-time sales and stock information. By providing access via a central API, the service is used for multiple purposes such as disposition or product availability information for consumers.

As stock calculation was only executed once a day for every store, a stock calculation failure would create a gap of stock-information for two days. Failed stock calculations resulted in wrong orders, out of stock situations in stores, and overstocking that led to write-offs and increased costs.

The problems compounded when in-stock store availability for promotions was not available for customers using the web or mobile application. This resulted in high call volumes on the customer hotline for stock-availability requests, and negative customer experience and customer churn if a certain product is not available in store.

By utilizing MongoDB's document model, the retailer's data schema was simplified and developer productivity increased. Without the need for JOINS, the retailer increased the stock-calculation speed to near real-time, and MongoDB analytics nodes allowed the retailer to separate productive calculation workloads from customer requests. Because of this, the calculation is stable and even special events, like holiday shopping season, have no impact on the stability and performance of the system.

Legal Notice

This document includes certain "forward-looking statements" within the meaning of Section 27A of the Securities Act of 1933, as amended, or the Securities Act, and Section 21E of the Securities Exchange Act of 1934, as amended, including statements concerning our financial guidance for the first fiscal quarter and full year fiscal 2021; the anticipated impact of the coronavirus disease (COVID-19) outbreak on our future results of operations, our future growth and the potential of MongoDB Atlas; and our ability to transform the global database industry and to capitalize on our market opportunity. These forward-looking statements include, but are not limited to, plans, objectives, expectations and intentions and other statements contained in this press release that are not historical facts and statements identified by words such as "anticipate," "believe," "continue," "could," "estimate," "expect," "intend," "may," "plan," "project," "will," "would" or the negative or plural of these words or similar expressions or variations. These forward-looking statements reflect our current views about our plans, intentions, expectations, strategies and prospects, which are based on the information currently available to us and on assumptions we have made. Although we believe that our plans, intentions, expectations, strategies and prospects as reflected in or suggested by those forward-looking statements are reasonable, we can give no assurance that the plans, intentions, expectations or strategies will be attained or achieved. Furthermore, actual results may differ materially from those described in the forward-looking statements and are subject to a variety of assumptions, uncertainties, risks and factors that are beyond our control including, without limitation: our limited operating history; our history of losses; failure of our database platform to satisfy customer demands; the effects of increased competition; our investments in new products and our ability to introduce new features, services or enhancements; our ability to effectively expand our sales and marketing organization; our ability to continue to build and maintain credibility with the developer community; our ability to add new customers or increase sales to our existing customers; our ability to maintain, protect, enforce and enhance our intellectual property; the growth and expansion of the market for database products and our ability to penetrate that market; our ability to integrate acquired businesses and technologies successfully or achieve the expected benefits of such acquisitions; our ability to maintain the security of our software and adequately address privacy concerns; our ability to manage our growth effectively and successfully recruit and retain additional highly-qualified personnel; the price volatility of our common stock; the financial impacts of the coronavirus disease (COVID-19) outbreak on our customers, our potential customers, the global financial markets and our business and future results of operations; the impact that the precautions we have taken in our business relative to the coronavirus disease (COVID-19) outbreak may have on our business and those risks detailed from time-to-time under the caption "Risk Factors" and elsewhere in our Securities and Exchange Commission ("SEC") filings and reports, including our Quarterly Report on Form 10-Q filed on December 10, 2019, as well as future filings and reports by us. Except as required by law, we undertake no duty or obligation to update any forward-looking statements contained in this release as a result of new information, future events, changes in expectations or otherwise.



The Next Steps to Modernization

Now that you've made it through the chapters of our overview, take the next step to modernizing with MongoDB. We work hand-in-hand with retailers to build customized solutions for the most important use cases improving the industry today. As a refresher, here's a brief outline of our most pertinent retail solutions:

1. Building a Modern Product Catalog
2. Application-Driven Analytics in Retail
3. Personalization and Search
4. The Single View: Constructing Customer 360
5. Getting Loyalty Programs Right with Receipt Mining
6. Achieving Greener Commerce
7. Enabling Mobile Solutions in the Retail Value Chain

Learn more about MongoDB for retail at our [solution homepage](#), or [contact us](#) to get the conversation started around what MongoDB Atlas can do for you.

Resources

For more information, please visit mongodb.com or contact us at sales@mongodb.com.

