

# Serving the Digitally-Oriented Consumer

Reimagining the Retail Experience with MongoDB  
August 2015

# Table of Contents

- Introduction 1
- A New Approach to Engage the Modern Shopper 1
- New Retail Challenges 2
  - Omni-channel View of Business 2
  - Customer Single (360 degree) View 2
  - Real-Time Inventory & Supply Chain 2
  - Developing an Innovative eCommerce Platform 3
  - Optimizing the Mobile Experience 3
  - Take Advantage of Social 3
- What is MongoDB? 3
  - Dynamic Schema 3
  - Querying 4
  - Sharding 4
  - Replication 4
  - Security 4
  - Ease of Development 4
- Business Value of MongoDB 4
  - Provide the Optimal User Experience 4
  - Rapid innovation with agile development 5
  - Lowered Total Cost of Ownership 5
- Customer Case Studies 5
  - Otto 6
  - CustomInk 6
  - Under Armour 6
  - eBay 7
- Conclusion 7
- We Can Help 7

## Introduction

Today's digitally-oriented society has created a new generation of shopper – one who lives across channels and expects personalization, perfect information, and instant gratification. These evolving consumer behaviors and expectations can be an incredible challenge for retailers, but they also represent a tremendous opportunity for leading brands to further differentiate themselves from the competition with technologies and processes that enable the omni-channel retail experience.

## A New Approach to Engage the Modern Shopper

Increasingly unpredictable customer demand is causing many retailers to intelligently adjust their backend systems. With the rise in the number of customer-facing channels, retail is no longer a one-way street from the store to the buyer. Consumers purchase items online and on their mobile phones; they expect to be able to schedule delivery to the location of their choosing or pick up products at the nearest physical store. Social media is also changing the

buying process; for example, demand for products can spike or plummet due to unforeseeable viral content.

Leading retail organizations have to be ready to handle all of these scenarios and more. This means not only ensuring that their customer-facing properties deliver the optimal experience, but also that their inventory management and supply chain management systems are integrated and updated continuously in a sustainable fashion. There is nothing worse for a business than to turn away a potential buyer because the item she wants is out of stock at the store and the retailer has no alternative means of procuring it for her. It can be equally as damaging to the business to have shelves stocked full of items that are not moving at all. Retailers need to find ways to meet customer demand in real-time in this increasingly complex landscape using cost-effective methods.

Processes and technologies that enable retailers to more accurately predict and address cross-channel demand are invaluable. Unfortunately, they are very difficult to implement on legacy systems. Siloed pools of information collecting in disparate data stores make it nearly impossible for retailers to develop a system-wide,

up-to-date view of their business. To start, retailers need regularly updated cross-channel views of their inventory to allow them to more accurately predict and address cross-channel product shortage and excess.

Consumers also demand and expect tailored, innovative shopping experiences. To develop these experiences, retailers need to make use of the oceans of multi-structured data that their customers produce. Unfortunately, rigid relational data stores make it cumbersome for project teams to incorporate new data types – like pictures, videos, geospatial information and other machine-to-machine sensor-driven data such as iBeacon, electronic shelf labels and RFID, etc.– into their offerings. Leading retailers need flexible, high performance systems that will not only scale seamlessly, but also allow them to do interesting things with their data. It starts with the right ecosystem, with flexible infrastructure where polymorphic, digital data can flow freely in the database.

## New Retail Challenges

### Omni-channel View of Business

Retailers have a common goal: to provide a seamless, branded experience across all selling and interaction channels. The growing number of backend source systems that need to be integrated and the increasing number and variety of customer-facing channels (smartphones, tablets, social, etc) make this singular consumer experience difficult to deliver with conventional tools. For example, retailers absolutely need to display consistent product availability information. Unfortunately, that can be difficult to deliver if their different channels receive information from different data stores.

MongoDB's document data model and dynamic schema make it ideal for combining information from different systems, allowing retailers to access a full, up-to-date view of their business (product catalog, inventory, etc). Aggregate information about customers' past interactions can also be stored in documents to provide a holistic, single view of the customer.

### Customer Single (360 degree) View

MongoDB's flexibility enables retailers to deliver a personalized customer-facing experience. For example, a customer may begin her shopping process by browsing online. Her clicks, likes, purchases, and other salient behaviors are recorded and ingested into MongoDB to construct her user profile for better content personalization and recommendations in her future interactions with the brand. In this example, her user profile represents a single view of her that the retailer can use to foster engagement and brand loyalty. With MongoDB, the organization can also aggregate all the information about her from their disparate data sources into a single document to gain a complete picture. She may "like" a particular item on the website, then continue the buying process on the organization's mobile application, which is also updated and personalized using the company's single view of her. Using the mobile application, she can either purchase the item outright or check for its availability at physical stores nearby and reserve it. Again, MongoDB's fully dynamic schema and flexible document data model makes it the ideal database for finding and delivering real-time product availability information fed from multiple sources. Finally, if she reserves the item for in-store pickup from her mobile app, sales agents in the store may receive the appropriate notification to hold it for her until she arrives.

### Real-Time Inventory & Supply Chain

MongoDB can also help ensure that all customer-facing properties are supported with the appropriate adjustments to inventory and supply chain management systems. In the example above, if the customer reserves the last item, a set of actions should take place to ensure that the store will be able to serve the next customer looking for the same item. This may involve a shipment from the nearest supplier or the nearest store with an excess in inventory, depending on which option is more cost-effective. Similarly, if the customer orders the item online, the organization should be able to parse through their fulfillment options and select the least costly method while still satisfying the customer's expected delivery time frame. By using MongoDB to develop a holistic view of their businesses, retail organizations can more easily build the operational requirements needed to support their omni-channel initiatives.

MongoDB's support for ad hoc analytics also allows retailers to implement intraday (or even hourly) inventory updates to help them identify and address trends faster.

## Developing an Innovative eCommerce Platform

Today, fast and responsive websites are table stakes, and consumers expect much more. Shoppers have grown accustomed to eCommerce features such as rich search, full product information with real-time availability, personalized content (recommendations, landing pages, etc) based on stored customer profiles, complete views of past interactions with brands, and more. MongoDB provides the database backbone to enable innovative eCommerce. Its powerful query functionality and indexing allows organizations to run analytics in real time against highly variable, fast-moving data as it is generated across any channel. MongoDB's native Aggregation Framework and MapReduce features allow retailers to do more complex real-time data processing. When used in combination with an analytical data store like Hadoop for batch and interactive data processing jobs, MongoDB can act as an enriched operational database that continually improves.

MongoDB's scale out architecture and native replication ensures high performance and high availability, allowing retailers to deliver and track their interactions with millions of customers. Furthermore, MongoDB's inherent flexibility and developer friendliness give retailers the ability to iterate quickly and accelerate their time to market for new eCommerce features.

## Optimizing the Mobile Experience

Mobile devices provide a storefront that's always open, allowing consumers to browse, favorite, save, and buy at all hours from any location. For retailers that want to be wherever their customers are, delivering an optimized mobile experience is key to staying competitive in a landscape that sees disruptive entrants every year. For example, mobile platforms can be used to serve up personalized offers that incorporate customers' geolocation data, stored customer profiles, and real-time inventory at stores in their vicinity. Retailers are turning to MongoDB,

with its native support for geospatial data, to help make these highly tailored consumer experiences a reality.

## Take Advantage of Social

Consumer-produced content is changing retail. No longer are conversations dictated by the retailer. Shoppers are talking back in the form of comments, tweets, likes, pins, reviews and videos. Shoppers are now also speaking with one another to compare notes and share recommendations or warnings. Retailers should not only be listening in, but also monitoring and responding to these conversations to try and create informed action items. They need to capitalize on the positive exchanges to take advantage of actionable insights that their internal teams missed, use social insights to enhance promotions, and be quick to respond to exchanges that paint them in a negative light.

To take advantage of this increasingly important social aspect of retail, organizations need a technologies that will allow them to ingest massive volumes of fast-moving data of all types. They need to be able to effectively address the increasingly large proportion of semi-structured, unstructured, and polymorphic data in today's user-generated information. Finally, they need to be able to efficiently test what works and what doesn't work in their attempts to formulate meaningful relationships between the data to inform their business decisions. MongoDB's inherent scalability and flexibility not only ensure that it can handle the volume, velocity, and variety of unstructured social data, but also that project teams using the database can quickly iterate in their efforts to make social data an integral part of their business.

## What is MongoDB?

Retail organizations are increasingly turning to MongoDB, the fastest growing database technology, to help them deliver on the omni-channel imperative. MongoDB is an open-source database used by companies of all sizes, across all industries and for a wide variety of applications.

## Dynamic Schema

While traditional databases store data in rows and columns, MongoDB stores data in the form of JSON documents.

Documents consist of one or more fields, with each field containing values of a specific data type, including integers, arrays, or even sub-documents. Documents that share a similar structure are organized as collections. Collections in MongoDB are analogous to a table in a relational database, documents are similar to rows, and fields can be thought of as columns.

Whereas relational databases require predefined structures and fixed relationships, otherwise known as a fixed schema, MongoDB's data model is inherently more flexible. Fields can vary from document to document; there is no need to declare the structure of a document beforehand. Fields can be added to a document without affecting any other documents in the system and without taking the system offline. MongoDB's dynamic schema allow organizations to seamlessly add and remove new data types as they please, allowing them to test and iterate on their ideas faster.

## Querying

Found in most databases, indexes are a crucial mechanism for optimizing system performance. MongoDB supports many types of indexes on any field in a document. MongoDB also supports powerful query functionality, which can return a document or a subset of specific fields within a document. When used in concert, these capabilities allow organizations using MongoDB to run ad-hoc analytics to extract real-time intelligence from their pools of multi-structured data.

## Sharding

MongoDB scales horizontally through a process called sharding, in which data is distributed across multiple physical partitions (known as shards). Sharding allows organizations deploying MongoDB to address the hardware limitations of a single server, such as bottlenecks in RAM or disk I/O, without adding complexity to the application. Distributing data across multiple shards allows organizations to easily address increasing data volume and throughput. Organizations can scale their MongoDB

deployments across fleets of commodity servers either on-premises or in the cloud, ensuring lower TCO as they grow their product catalog and user base.

## Replication

MongoDB stores multiple copies of data in the form of replica sets through a natively supported process called replication. Replica sets can be thought of as fully self-healing shards that automatically ensure high availability; failover happens without the need for an administrator to intervene.

A replica set consists of a configurable number of replicas; the higher the number, the greater the protection against database downtime. To ensure consistency by default, one member of the set acts as the primary replica for reads and writes while the other members act as secondaries. If the primary member fails for any reason, one of the secondary members is automatically elected as the new primary and begins to accept all writes to the database.

## Security

MongoDB's native enterprise-grade security features such as Kerberos authentication and SSL encryption ensure that retail organizations can build PCI DSS compliant systems. Other security features include collection-level authorization, user-defined roles and field level redaction. MongoDB Enterprise customers can also take advantage of auditing, LDAP integration, and x509 authentication.

## Ease of Development

MongoDB supports a rich library of idiomatic drivers for many of the most popular programming languages, including C, C++, Java, Perl, Python, Ruby and more. With MongoDB, development teams familiar with object-oriented programming languages also do not have to worry about managing the object-relational impedance mismatch as their data is already modeled like an object.

# Business Value of MongoDB

## Provide the Optimal User Experience

It is critical for companies in retail to deliver an optimized user experience to their customers. The difference between keeping a shopper loyal to the brand and losing them forever can be a single bad customer experience. Personalized cross-channel shopping experiences are quickly becoming the norm. To grow and retain their customer base, retail organizations need systems that will not only allow them to build innovative front-end experiences, but also support the backend processes that make those experiences possible.

## Rapid innovation with agile development

Organizations need solutions that are easy to implement; technologies with steep learning curves can bring teams to a complete standstill. Database technologies like MongoDB, which include a rich library of idiomatic drivers for the most popular programming languages, ensure that project teams will have a much easier time being productive.

To adapt to evolving business requirements, organizations need flexible, iterative development practices that make it easy for their project teams to adjust their plans on the fly. The underlying database technology must be also flexible enough to accommodate these adjustments without hampering developer productivity.

## Lowered Total Cost of Ownership

Infrastructure and operating costs have fallen dramatically in recent years, driven mostly by the increasing popularity of low-cost commodity servers and on-demand cloud computing services. However, traditional database technologies, most of which are designed to scale vertically in a single datacenter, are not suited to take advantage of these developments. Retail organizations that want to take full economic advantage of commodity servers and the many benefits of cloud computing should turn instead towards database technologies that natively scale out horizontally.

# Customer Case Studies

## Otto



Otto's otto.de website is the largest web property for women's and children's clothing in Europe. The company found that they were having difficulty providing differentiated, personalized offerings to their millions of website visitors using conventional tools. They needed an alternative approach that would allow them to quickly iterate and implement their ideas, provide real-time information on the 2 million plus products in their catalog, and deliver personalized experiences to each of their 30 million unique users. After extensive trials, Otto chose MongoDB to support their next generation eCommerce platform on otto.de. MongoDB allows Otto to personalize content for their audience and deliver dynamic product views and web pages.

Otto was able to aggregate data from their different source systems using MongoDB's dynamic schema and flexible document data model. This allowed the company to easily construct and share consistent views of their product catalog. MongoDB's scale-out architecture allowed Otto to meet their low latency requirements and ensure high performance. Otto was also able to achieve high availability using MongoDB's native data replication.

"With MongoDB, we chose a partner who could really support us in this process with MongoDB consultants helping us in both design and training ... As a result, we have a modern, digitally-oriented application development environment which will allow us to implement our innovative ideas as quickly as we create them."

## CustomInk



Founded in 2000, CustomInk.com allows their customers to design and order custom-designed products such as t-shirts and embroidered caps. The company was named in Inc. Magazine's Inc. 500 list of America's fastest-growing private companies.

CustomInk turned to MongoDB to evolve their eCommerce business when they found that their Oracle database limited their ability to deliver new web functionality quickly. Their decade-old product catalog was restricted by a rigid, static data structure that made it difficult to create new merchandising opportunities. Finally, their reliance on Oracle limited them to one datacenter, limiting their disaster recovery options and preventing them from moving their applications into the cloud.

With MongoDB, CustomInk now deploys their applications wherever they want - locally, in a datacenter, or in the cloud. The company currently uses MongoDB for three use cases:

'Single View' Logging - By aggregating all application and web site activity data on MongoDB, CustomInk eliminated the need to track transactions across multiple servers. CustomInk can now easily capture metrics in order to better understand customer behaviors.

Content Management – MongoDB's document model and dynamic schema enabled rapid development of CustomInk's next-generation product catalog and provided more dense data storage. They were able to easily model product data and eliminate fields with empty values. A baseball cap, for example, should not be represented in the same structured way as T-shirts; CustomInk does not need

to associate multiple sizes with caps, nor do they need to know whether they can "print on the sleeve." MongoDB also provided flexibility to iterate over the design and update the application with new fields, a benefit that will continue to serve them well as requirements change.

Supply Chain Management – Before switching to MongoDB, CustomInk managed their complex supply chain with Excel spreadsheets shared across the organization. Today, all data related to the printer, vendor and supplier network – e.g., inventory data, ordering logistics, printer capabilities, location and availability – is stored in a MongoDB web application that enables and automates their entire supply chain. MongoDB directly impacted how CustomInk processes orders internally, and even led them to readjust several operations roles.

## Under Armour



Founded in 1996, Under Armour is a leading developer, marketer and distributor of branded performance apparel. Under Armour generates more than \$2 billion in annual sales through retailers and increasingly through their own online shop. In 2011, the company decided that they needed to improve the usability of their online store and address limitations with their existing content management system. Due to high growth, the company was also running into site scalability issues with their existing platform built on SQL Server.

After evaluating a number of commercial eCommerce solutions, the company decided that building a custom solution gave them the best chance of releasing an upgraded website that would meet all their requirements by the 2011 holiday shopping season. Building a custom solution allowed them to reuse components of their existing platform, have direct control during implementation, and incorporate a high level of customization. With only eight months to design, test, and

deploy their new online store, Under Armour turned to MongoDB.

MongoDB's native replication and sharding capabilities addressed the company's high availability and scalability requirements. Under Armour's development team also found MongoDB's dynamic schema to be particularly useful during implementation as it allowed them to rapidly change the store to meet the changing demands of the business. Finally, MongoDB's high performance on both reads and writes guaranteed that site visitors experienced an eCommerce system that was incredibly fast.

"The flexibility of MongoDB has put a lot of power in the hands of our developers. MongoDB provides us a lot more flexibility in how we think about our database. This is particularly true with our page model. We can define each page of the website in the data rather than in the code, while our runtime engine parses the page definition to determine which layout, which modules, which content, and which data should be rendered on the page. And all of this is controlled seamlessly from a relatively simple Flex UI."

— Brian Massey, Director of eCommerce Information Systems, Under Armour

Under Armour currently stores the majority of their data in MongoDB, including everything from Adobe Scene7 image recipes to page definitions and product data. The company is currently planning out the technical architecture for the next version of Under Armour website, in which MongoDB will once again play a prominent role.

eBay



Founded in 1995, eBay is a multi-billion dollar business specializing in consumer-to-consumer transactions. Their existing systems, based on Oracle, already scales to thousands of developers, petabytes of data, and hundreds of millions of SQL queries per day, but eBay looked to MongoDB to improve their stack and experiment with new ideas. They sought out a solution that would allow them to:

- Achieve faster development cycles
- Add new attributes faster and without the involvement of a database administrator
- Scale quickly
- Deploy new databases to the cloud on demand

MongoDB allowed eBay to meet all these criteria. eBay's developers cited MongoDB's flexible document model, easy scaling (via on demand sharding), and compatibility with cloud hosting providers as key features.

MongoDB is currently used in several use cases. These include:

Media Metadata – MongoDB stores all media metadata for eBay's website. This mission-critical application includes references to pictures of every item for sale on eBay.

Search Suggestion – MongoDB is used to hasten search suggestion delivery. To ensure an optimal customer experience, search suggestions must be served within 60-70 milliseconds. By storing the search suggestion list in a MongoDB document, the system can run a query in less than 1.4 milliseconds.

Merchandising Information Cache – MongoDB powers the merchandising backend, which eBay uses for product classification and categorization. MongoDB's dynamic schema proved essential as new attributes and new queries are constantly being added to the system.

## Conclusion

The demanding expectations of the digitally-oriented consumer are forcing retailers to reimagine the retail experience. They are realizing that to remain relevant in this highly crowded and competitive global market, they need to deliver up-to-date product catalogs, personalized

recommendations, timely offers and more, all on a cross channel platform to foster brand loyalty. To support these omni-channel goals, retail organizations need to install new backend processes to intelligently align their inventory and supply chains to closely match increasingly unpredictable consumer demand. All of these initiatives and more are made easier with a flexible database like MongoDB.

## We Can Help

We are the MongoDB experts. Over 2,000 organizations rely on our commercial products, including startups and more than a third of the Fortune 100. We offer software and services to make your life easier:

**MongoDB Enterprise Advanced** is the best way to run MongoDB in your data center. It's a finely-tuned package of advanced software, support, certifications, and other services designed for the way you do business.

**MongoDB Cloud Manager** is the easiest way to run MongoDB in the cloud. It makes MongoDB the system you worry about the least and like managing the most.

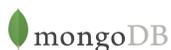
**MongoDB Professional** helps you manage your deployment and keep it running smoothly. It includes support from MongoDB engineers, as well as access to MongoDB Cloud Manager.

**Development Support** helps you get up and running quickly. It gives you a complete package of software and services for the early stages of your project.

**MongoDB Consulting** packages get you to production faster, help you tune performance in production, help you scale, and free you up to focus on your next release.

**MongoDB Training** helps you become a MongoDB expert, from design to operating mission-critical systems at scale. Whether you're a developer, DBA, or architect, we can make you better at MongoDB.

Contact us to learn more, or visit [mongodb.com](http://mongodb.com).



New York • Palo Alto • Washington, D.C. • London • Dublin • Barcelona • Sydney • Tel Aviv  
US 866-237-8815 • INTL +1-650-440-4474 • [info@mongodb.com](mailto:info@mongodb.com)  
© 2015 MongoDB, Inc. All rights reserved.

## Resources

For more information, please visit [mongodb.com](http://mongodb.com) or contact us at [sales@mongodb.com](mailto:sales@mongodb.com).

Case Studies ([mongodb.com/customers](http://mongodb.com/customers))

Presentations ([mongodb.com/presentations](http://mongodb.com/presentations))

Free Online Training ([university.mongodb.com](http://university.mongodb.com))

Webinars and Events ([mongodb.com/events](http://mongodb.com/events))

Documentation ([docs.mongodb.org](http://docs.mongodb.org))

MongoDB Enterprise Download ([mongodb.com/download](http://mongodb.com/download))