


# Top Five Global Retailer Modernizes Their Online Shopping Experience

YugabyteDB Underpins A Modern Shopping Experience For Millions Of Customers and Third Party Sellers


Retail




**Company**  
Top 5 Global Retailer

**Location**  
US


### Main Use Cases




Database Modernization



Cloud Native Microservices



Geo-Distributed Apps



Edge / IoT

## Executive Summary

As one of the world’s largest retailers, this company is focused on constant innovation to maintain its competitive edge. Their innovation focus accelerated when they saw the opportunities possible with new cloud technologies. But, as first movers often discover, being a leader means you don’t have the benefit of learning from others.

First generation cloud databases offered the retailer the scale and resilience they needed to meet their anytime, anywhere, online shopping vision. But, they lacked transactional consistency, which made many customer and vendor use cases hugely complex to implement and maintain. To fulfill their modernization goals and overcome this obstacle they started looking at database solutions, which started their journey to distributed SQL.

Through discussions with Yugabyte, they discovered that [YugabyteDB](#) could deliver the unlimited scale and resilience of a cloud native database, with the familiarity of SQL, plus transactional consistency. Their use of [distributed SQL](#) has led to massive savings as a result of accurate data and high availability across cloud regions.

## About The Customer

This top five global retailer has thousands of stores in the US and internationally. By embracing innovation, they have created a seamless experience that allows customers to shop anytime and anywhere, online and in stores. They are creating opportunities and bringing value to customers and communities around the globe. Their eCommerce website lists millions of items and they are one of the largest employees in the world.

## Key Results



**100M+**  
items in the product catalog



**\$10M**  
in lost revenue recovered



**>250K**  
queries per second

## The Challenge

This global retail giant operates a huge e-commerce site featuring thousands of sellers and millions of products. Underpinning this offering is a number of diverse applications, handling large volumes of data. These applications manage how products are listed, displayed, grouped, filtered, searched, and presented to give the user a personalized but seamless buying experience. Other applications deliver a simple, effective ‘shop-window’ to marketplace suppliers.

These wide range of applications have been architected as a portfolio of microservices, powered by a cloud native technology stack exploiting Kafka, Akka Streams, and using Apache Cassandra as the core database. Services are deployed in a multi-datacenter topology, using a multi-cloud deployment strategy.

The technology team that builds and runs the retailer’s site platform approached Yugabyte looking for an innovative and reliable technology solution. They were having major issues with consistency in the product catalog. This requires high velocity, high volume, streaming updates, from millions of market vendors to be reflected in the Product Catalog as they arrive—and apply consistently across the network.

Their current Cassandra implementation was unable to guarantee this consistency. The cost of this was considerable and unacceptable. It resulted in user dissatisfaction, and an annual loss of millions of dollars. The loss was a result of having to fulfill orders promised to customers from expensive third party vendors because inaccurate inventory numbers between regions did not show the item as sold out in a timely manner.

Another key challenge was how to ensure a truly always-on business in the event of a major cloud failure or outage. Their existing environment did not offer the ability to seamlessly keep their key applications running with the latest data in the event of a major region outage—something that would actually happen to them soon after our initial conversations.

Their IT team had the skills, experience, and technology stack in place to deliver cloud native applications rapidly. But the level of consistency-management in their Cassandra database was inadequate for the systems of record that formed the core of their responsibilities.

## Database Requirements

- Deliver strong consistency
- Deliver scale and performance
- Offer high availability and strong resilience
- Support different datacenter topologies
- Support deployment to multi-cloud solutions

They evaluated Azure CosmosDB, Azure Cloud SQL, and a range of other databases, but none of these candidates were able to meet their requirements. When YugabyteDB was presented, the team realized it had the potential to meet all of their criteria and was swiftly identified as a prime candidate for a Proof of Concept (PoC) implementation.

As well as meeting their key requirements, YugabyteDB offered the key advantage of being 100% open source and built on open standards. This provided them with the confidence that they were not locked into a proprietary technology.

## Solution

The retailer’s evaluation team prioritized two critical use cases that must be implemented to address the inconsistency that plagued their existing platform and affected their users’ experience. The first was the subject of the PoC. The intention was to drive the use case into production as quickly as possible, then fast-track the second close behind it. This meant that the demands of the second use case were also a significant component of the evaluation.



### Use Case 1: Product Catalog Matching

**The matching service consolidates items for publication on the marketplace and groups them, based on multiple characteristics, to create a unique ID in the product catalog and across other services and applications.**

**The core requirements for this use case were:**

- Zero downtime
- Massive concurrency and parallelism
- Linear scaling
- High availability



## Use Case 2: Primary Key/Alternate Key (PK/AK)

- PK/AK is the key-to-value mapping service that supports lookups across various key identifiers in different services and systems.
- The workload pattern for this service was 90% Reads and 10% Writes. Very low latency for the Reads was essential.

Because Cassandra was the incumbent database, the implementation of both of these use cases benefited from YugabyteDB's powerful support for the [Cassandra Query Language API \(YCQL\)](#). YCQL minimizes disruption for existing Cassandra-native applications and reduces the learning curve for the Cassandra-savvy engineering teams.

Unlike Cassandra, the YugabyteDB solution offers powerful [ACID transactional capability](#) to deliver multi-table transactions with strong consistency in a multi-region deployment. This was key to the success of the primary matching use case. YugabyteDB's ability to consistently achieve single-digit milli-second reads at scale in a multi-region deployment was essential to success in the PK/AK use case.

YugabyteDB's ability to support both RDBMS-style transactions and a NoSQL programming model on the same datastore made all the difference for this customer. Being able to deliver consistent data across cloud regions also met their key requirement for high availability and resiliency. YugabyteDB offered flexibility deployment options, allowing the company to synchronously replicate data across multiple Azure regions, their public cloud of choice.

## Results & Benefits

By embracing distributed SQL through YugabyteDB, the retailer enjoyed a number of key technical results, including:

- **Very high-volume ACID-transactional consistency across a multi-region deployment on multi-table transactions.**

- **Consistent very low latency for high volume reads using 'Follower Reads' to deliver from the nearest replica to the user.**
- **Rigorous resiliency/failure scenario tests, including Region Failure Simulation, Cluster Scale Up, Rolling Version Upgrades; all of them successful with zero downtime to end-client applications.**
- **42 nodes deployed for the two use cases, across three regions.**

In addition to these technical wins, the overall business was able to enjoy additional benefits that they highlighted to their management team and leveraged to continue to expand their use of distributed SQL in the business. These key business results included:

- **Elimination of buyer and seller friction caused by Product Catalog inconsistencies.**
- **Elimination of time-consuming and error-prone remediation to restore consistency.**
- **Improved engineering capability to apply these benefits across a range of services.**

YugabyteDB also achieved another major win for the company when their always-on architecture was put to the test. In the US, this retailer operates across three Azure regions, with classic legacy databases for some applications, and YugabyteDB spread across three regions for other services.

When a major ice storm hit Texas in 2021, the retailer's applications powered by YugabyteDB had zero loss of data and traffic was automatically routed to other regions in just three seconds. This meant no action was required here by the IT team, who instead had to do manual reconciliation of their applications and databases on their US South-Texas cloud region (not powered by Yugabyte), which was offline for over four days.