

RETAILER CO-OP MIGRATES FROM DATA WAREHOUSE APPLIANCE TO AZURE SYNAPSE IN RECORD TIME

CUSTOMER IN BRIEF



COMPANY The Co-operative
LOCATION Manchester, UK

PLATFORM PARAMETERS

10+ TB
COMPRESSED
DATA

326,000
DAILY QUERY
REQUESTS

1,000+
BUSINESS
USERS

Co-operative Group Limited, also known as Co-op, is one of the world's largest consumer-owned co-operatives with more than 65,000 employees.

Co-op is a food retailer and wholesaler, the largest funerals provider in the UK, a major provider of regulated consumer legal services, particularly of probate and wills, and a major provider of life planning and insurance products.

STORY HIGHLIGHTS

Accelerated migration to cloud

Using Datometry Hyper-Q, Co-op moved its entire ETL and reporting system for its grocery business to Microsoft Azure Synapse in only 10 months.

Realized substantial cost savings

Avoiding a full rewrite of its codebase allowed Co-op to move to Azure at a fraction of the cost of a typical migration, without the risk of failure or budget runovers.

Migrated all users seamlessly

Co-op's migration with Datometry was so seamless that employees who weren't told about it had no idea it had even happened.

CHALLENGE

With more than 65,000 employees and revenue of nearly £11 billion pounds, the Co-operative Group has a data infrastructure to match the size of its business. Its grocery operation, whose 2,500 stores are familiar to anyone in the UK, is backed by a data warehouse organized in over 10,000 tables and ingests data from 885 scheduled loading processes. From this data, more than 6,000 reports are then served to about 1,000 field-based users.

Co-op regularly reviews the cost and plans for its systems, and in 2019 it set off to review its plans for its data warehouse appliance which would be out of support in 2021. The "defender" position was to stay with the existing vendor. As Co-op had built a datalake and data warehouse in Azure, its strategic "challenger" direction was to move to new cloud technology solution.

After a full review, taking into account external expert advice, and running a robust selection process, Co-op selected the cloud migration option. But while the benefits of cloud migration were obvious, they paled when considering the costs, time, and risk associated with doing so, which would involve a wholesale rewrite of their codebase, including their ETL pipelines and reporting systems. On top of that, Co-op was already in the midst of implementing a brand-new ERP system, requiring it to rewrite portions of its system; adding in a costly, time-consuming and risky database migration would have added multiples of complexity. "A conventional migration would have taken years" estimated Rob McKendrick, Head of Strategy and Engineering at Co-op.

SOLUTION

As a solution, Microsoft introduced Co-op to Datometry, a key ISV partner. The prospect of being able to move all of their existing applications, including their ETL setup, to Azure Synapse without having to execute a conventional database migration was immediately appealing to the Co-op team.

Migrating to Azure using a 5-step approach

The migration would be executed in five distinct phases. During the first phase, the Co-op team devised a complete migration plan together with Datometry Customer Success using the detailed analysis of their workloads from Datometry qInsight.

Next, they translated all existing data definition language (DDL) statements from the legacy system using Datometry qShift to generate the appropriate definitions natively in Synapse. This schema conversion comprised over 10,000 objects and saw fundamental tuning principles, such as specific table properties, applied automatically.

In phase three, the Co-op team designed and executed a migrating strategy for all historical data out of the

“We gave the first groups of users a heads up that we’re going to switch them over. However, after we saw how smooth the change was, we didn’t tell people that we’ll be replacing their core data platform and just let them know afterward.”

Rob McKendrick,
Head of Strategy and Engineering, Co-op

on-premises system and into the cloud. By managing data snapshots in a sophisticated way, they would later be able to replay workloads as needed and accelerate testing cycles greatly.

In the subsequent phase of testing, individual work-

loads were transferred to the new system. Functional testing of the primary ETL jobs was followed by testing of a subset of reports. Being able to include the business teams in the testing ensured IT had buy-in from the business.

Lastly, during the cutover phase, Co-op performed parallel runs for several weeks. During that time both systems, the legacy stack as well as the new Azure Synapse system would be loaded synchronously and offer the business an apple-to-apple comparison of the old and the new. More importantly, however, this approach facilitated a seamless cutover of individual user groups without any disruption to the business.

RESULTS

Migrating to Azure with Datometry, Co-op was able to perform a full migration from its legacy data warehouse in just 10 months with an incredibly small core team of only about a dozen people. Executing this migration the conventional way using transpilers and other tools to rewrite the entire corpus of SQL embedded in applications, scripts, and reports would have required involving all business users and bringing on additional IT staff for the duration of the migration. Even then, a conventional migration would have taken years and still run the risk of overrunning the given budget severely or failing entirely.

Transition with no interruption to the business

But perhaps the greatest testament to the migration’s success was how little of Co-op’s users actually knew it had even taken place. “It was literally a switch overnight, for all different types of users and all different types of data,” says McKendrick. “We gave the first groups of users a heads up that we’re going to switch them over. However, after we saw how smooth the change was, we didn’t tell people that we’ll be replacing their core data platform and just let them know afterward.”

ABOUT DATOMETRY

Datometry is the database virtualization platform for the cloud. Datometry Hyper-Q brings unprecedented simplicity to enterprises that want the benefits of modern cloud databases—without the complexity of reinventing existing applications and business processes. Find out more at www.datometry.com