

SNOWFLAKE YIELDS BETTER INSIGHTS AND EASES GDPR COMPLIANCE FOR SAINSBURY'S

RETAIL

Sainsbury's

COMPANY Sainsbury's
LOCATION London, United Kingdom

SNOWFLAKE WORKLOADS USED



Sainsbury's is the UK's second largest retailer with over 1,400 stores and a large digital presence. Its mission is to help customers live well for less. The company offers customers distinctive, quality products at competitive prices across food, general merchandise, clothing, and financial services products, including the brands Argos, Tu, Sainsbury's Bank, and Habitat. Sainsbury's multibrand, multichannel business provides choice, flexibility, and convenience for its customers.

STORY HIGHLIGHTS:

Near-zero maintenance

Snowflake helped Sainsbury's consolidate data from disparate on-premises systems without the need for infrastructure support.

Data science optimization

Snowflake's data science workload enabled Sainsbury's to build a new product-matching feature that customers love.

Time Travel

Snowflake Time Travel enabled Sainsbury's to apply iterative software engineering principles to its data processing.

“We took the financial heart of our business and moved all of the data processing from its original data store to Snowflake.”

—GRAEME FORBES, Principal Data Engineer, Sainsbury's

CHALLENGE:

Meeting customer needs and GDPR requirements with siloed data

Big plans were in the works at Sainsbury's headquarters. According to Sainsbury's Principal Data Engineer, Graeme Forbes, “It was decided that we would try to merge all of our different offerings into one multibrand, multichannel business and move from a model where data was localized, siloed, and duplicated to one where data is democratized and available to everyone.” In addition, Sainsbury's wanted to know its customers better than anyone else, understand the cost of its decisions, and remove barriers to innovation and curiosity.

Around this time, the EU rolled out the General Data Protection Regulation (GDPR), a set of data protection and privacy laws. Complying with GDPR and responding to customers' information requests presented a new set of challenges to Sainsbury's.

According to Graeme, “When you have data scattered throughout a business, data analysis is a very intensive, time-consuming process. It should be quick to find all the information about a particular customer, but with many different businesses, it's a hard job that requires a very large team.”

This pain point caused Graeme to explore ways to consolidate data and make compliance requests and audits easier to support. Instead of searching in seven or eight different systems, he wanted to consolidate data into a single system, one that would serve as the single source of truth.

1,000

Sainsbury's data analysts using Snowflake to perform queries and run reports more effectively

8:1

Reduction in systems, while providing enhanced GDPR support

One

Source of truth across brands and channels for data analysis

SOLUTION:

Data consolidation with Snowflake's Data Cloud

Sainsbury's chose the Snowflake Data Cloud to create a single source of truth for business-critical data across all of its brands. According to Graeme, "We took all of our data, remodeled it, and put it into Snowflake as our main consumption point." Sainsbury's uses Snowflake Streams and Tasks to process transaction stream data, as well as clickstream data from their websites.

Effective data mapping with Data Vault

Sainsbury's implemented Data Vault, a method and approach to modeling an enterprise data warehouse that is agile, flexible, and scalable. Data Vault enabled Graeme to map ownership to the data model. According to Graeme, "I could say that this vault is all product data, that vault is all customer interaction, and another vault is all finance and commercial information."

Data Vault served as the hub where data was written. Satellite tables build off keys that are written to the hub. The satellite tables enhance the data. "You have the link tables that are the actual driving engines. They're the things that link data from multiple different satellites and hubs, and those link tables are then very traceable," said Graeme.

Data Vault lets Graeme see where the data is, where it came from, and how complex it is. In addition to enforcing ownership, Data Vault allows Sainsbury's to build a map that is business-focused rather than data-focused.

Graeme said, "In the past, our data vault had performance limitations, since complex joins happened on the same resource where data was stored. With the separation of storage and compute in Snowflake, that limitation goes away."

"Snowflake enables us to complete complex queries without disrupting our production data. In addition, per-second pricing tells us exactly how much it costs."

—GRAEME FORBES, Principal Data Engineer, Sainsbury's

Iterative data processing with Time Travel

Snowflake Time Travel enables accessing historical data (data that has been changed or deleted) at any point within a defined period. Time Travel helped Sainsbury's apply software engineering principles to its data processing, iterating until it found the approach that worked best. The iteration was made possible because Time Travel can clone databases at particular points in time at a granularity of minutes or even seconds.

According to Graeme, "We can run our processing until it works or make lots of little changes until it's finally ready. Before Snowflake, we needed to spin up and manage a secondary database. On Snowflake, we can do this in parallel with the production system."

RESULTS:

New benefits and services based on data science

With Snowflake, Sainsbury's was able to launch a product-matching service that compares its products with competitors' products. The service is a popular feature on its websites. According to Graeme, "The capability uses a data science platform that's enabled by data being in one place in Snowflake."

Graeme and his team built business models to calculate the business benefits provided by the data science engine on Snowflake. "We evaluate where we were and the methods we previously used and compare that to the predictive output. We process different events and logging methods and create a real-time dashboard that shows benefits like time savings and accuracy."

Easier GDPR compliance

Before Snowflake, handling GDPR requests and generating reports entailed complex processes on multiple systems, including IBM MDM, Oracle, and Postgres. Graeme and his team often couldn't trace the authoritative source for a particular data type.

Snowflake as the single source of truth makes it easier for everyone.

"Snowflake helped us eliminate layers of duplication and confusion. We're no longer looking in seven or eight different systems. We're now just looking in one."

—GRAEME FORBES, Principal Data Engineer, Sainsbury's

FUTURE:

Team ownership of end-to-end solution

Sainsbury's has grand ambitions. Graeme wants to use technologies like Snowflake to become a software engineering house and create full ownership of systems end-to-end. Snowflake will help enable Sainsbury's move from a model of separate teams owning parts of the solution to one where teams own the end-to-end solution, from operational support to governance to customer experience. With Snowflake, Sainsbury's will deliver high-quality, long-lasting products to customers in an even more efficient manner.

ABOUT SNOWFLAKE

Snowflake delivers the Data Cloud—a global network where thousands of organizations mobilize data with near-unlimited scale, concurrency, and performance. Inside the Data Cloud, organizations unite their siloed data, easily discover and securely share governed data, and execute diverse analytic workloads. Wherever data or users live, Snowflake delivers a single and seamless experience across multiple public clouds. Join Snowflake customers, partners, and data providers already taking their businesses to new frontiers in the Data Cloud. [snowflake.com](https://www.snowflake.com)