



BREAKING DATA SILOS:

How 8 companies gained greater data warehousing value with Starburst

The modern data lake analytics platform

Contents

BestSecret optimizes analytics and reduces data warehousing costs	7
Comcast speeds up analytics with a hybrid cloud data architecture	8
El Toro accelerates time-to-market for advertising campaigns	9
Domino Data Lab delivers access to data exploration	10
doxo enhances agility, materialized views, and database analytics	11
Online food ordering leader improves query performance 10-15x	12
Biopharmaceutical company speeds time-to-insight for faster drug development	13
Food & beverage giant accelerates cloud migration	14
Complement your data warehouse with Starburst	15



Today's companies face significant challenges when it comes to managing their data warehouses, architectural complexity, data silos, and the need for a single source of truth.

One of the common challenges faced by IT teams is the difficulty of accessing data residing on multiple platforms. This has very real consequences ranging from lost productivity to incomplete datasets — which leads to slower, less accurate trend analysis. IT and business leaders are looking for solutions that enable them to easily access and analyze data from diverse sources, resulting in faster time to insight and reduced total cost of ownership (TCO).



This is where Starburst enters the picture.

Starburst runs high performance data lake analytics directly on object storage and helps you choose what type of data should move into the data warehouse. Because modern data lakes can give you warehouse-like capabilities, you don't need to put all your data into a single data warehouse. For organizations relying on a data warehouse to centralize data for all analytics, Starburst can help you de-risk that investment and make the warehouse a component of a broader modern data lake strategy.

Companies that partner with Starburst see results, such as:

100%

future-proof
architecture

53%

lower TCO

90%

faster time
to insight

In this eBook, we'll look at eight companies that saw these results and more by complementing their data warehouse with Starburst.



BESTSECRET

European retailer BestSecret optimizes analytics and reduces data warehousing costs

Overview

BestSecret is a leading European members-only online destination for premium and luxury off-price fashion. By harnessing data-driven decision-making, it optimizes product offerings and inventory management, providing its members with a seamless and delightful shopping journey.



"We moved from a monolithic Snowflake approach to a decentralized approach with Starburst and Iceberg. Now we can skip the data warehouse step completely, and complete analytics on the data right where it sits."



Lutz Künneke,
Director of Data Engineer

Challenge & Solution

BestSecret has been using Snowflake for three years. The engineering team implemented Open Source Trino for cost saving measures, as Snowflake was getting too expensive. Snowflake also couldn't connect to the rest of their data sources, delaying analytics.

The retailer deployed Starburst as the Enterprise-grade version of Trino. Starburst has the ability to federate BigQuery and Snowflake, completing the retailer's data stack.

What Starburst Delivered

- 70% cost reduction after moving workloads from Snowflake to Starburst
- Better, faster business decisions (specifically for monitoring campaign performance, pricing, and sales)
- Predictable costs & performance of data analysis



Comcast speeds up analytics with a hybrid cloud data architecture

Overview

This telecommunications company is the leading pay-TV and cable TV company in the United States. With more than 15 million subscribers, Comcast retains and ingests tremendous volumes of data across various platforms.



“When end users are going into on-prem or cloud environments, they will be presented with all the data sets they have access to, irrespective of where the data is located. This offers a huge value to our end users.”



Bryan Aller, Director of Software and Engineering

Challenge & Solution

Over time, Comcast’s growth led to increasingly complex and varied data storage infrastructure across the company in which data was siloed on different platforms – from RDBMS, data warehouses (Snowflake and Teradata), NoSQL databases, and data lakes. It took months to run ETL jobs, hindering its ability to make quick business decisions.

To solve the data bottlenecks, Comcast transitioned to a cost-effective hybrid cloud data architecture with Starburst as the query engine to federate data access.

What Starburst Delivered


- 250-300 TBs data pulled daily into the platform
- Real-time data exploration and results
- Reduced enterprise data warehouse spend, ETL, and labor costs



El Toro accelerates time-to-market for advertising campaigns

Overview

Since its inception in 2014, advertising technology company El Toro has become a leader in IP targeting. Its patented technology allows advertisers to target individual homes based on IP addresses and the demographic information of the people living there with extreme precision.



“We evaluated Snowflake, but given the incredibly ad-hoc nature of our business it wouldn’t be cost effective. We would have to increase our cost by 10X to achieve the performance that Starburst offers us at a fraction of the cost.”



Richard Teachout,
Chief Technology Officer

Challenge & Solution

El Toro's customers required greater data-driven insights from their IP targeting platform, and moving to a data warehouse wasn't a scalable option. Adding a data warehouse to the stack would add complexity through data copying, and storage costs would've escalated to millions of dollars.

The company built a data lake architecture with Starburst Enterprise in order to run fast, live, interactive queries on its data lake. Today, massive amounts of structured and unstructured data is available faster to end users, leading to an accelerated time-to-market

What Starburst Delivered


- 300% improvement in query performance
- \$5 million in savings over the cloud data warehouse
- Thousands of campaigns launched every day with quick, easy data access



Domino Data Lab delivers access to data exploration

Overview

Domino Data Lab powers model-driven businesses. Its leading Enterprise Machine Learning Operations (MLOps) platform accelerates the development and deployment of data science work while increasing collaboration and governance.



“The data warehouse and data lake space is constantly evolving, and our enterprise focus means we have to support customer requirements across different platforms. Starburst gives us the ability to move quickly to support ever-changing use cases within complex enterprise environments.”



David Schulman,
Head of Partner Marketing

Challenge & Solution

Domino was challenged to enable secure and seamless data access to the plethora of data sources its customers worked with.

After evaluating the option to build data connectors in-house, the company decided to partner with Starburst to ensure data sources are easily accessible and shareable. Starburst’s data lake analytics platform provides SQL-based access to data, unlocking data for Domino’s customers at a fraction of the cost of cloud data warehouses.

What Starburst Delivered

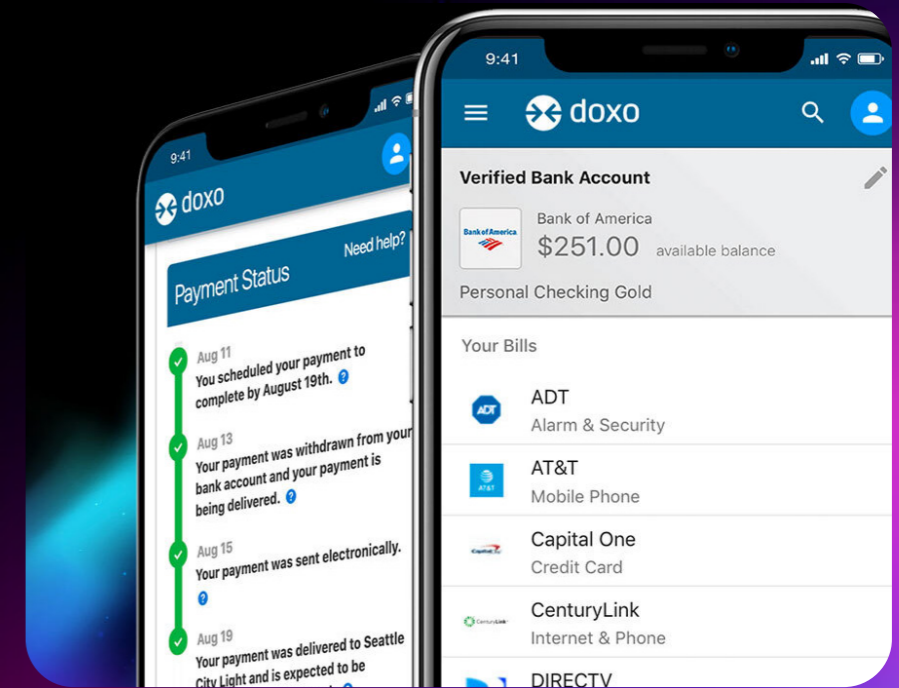
- 90% reduction in connector development time
- 3X cost savings
- Improved customer satisfaction



doxo enhances agility, materialized views, and database analytics

Overview

A personal finance company, doxo helps more than six million people protect their financial health and accomplish their financial goals by breaking free of the bill-paying burden. The company facilitates millions of transactions daily, allowing people to manage their due dates and pay all their bills in one place.



“Starburst allows us to satisfy all of our data consumers at doxo with the least amount of effort on the part of the people who manage the data assets.”



Daniel Cook, Sr. Product Director,
Data Platforms and Analytics

Challenge & Solution

doxo had data siloed in different formats on various platforms, including Redshift, MySQL, Snowflake, Hive, and Amazon S3. Generating daily insights was critical to the health of the business, but the process of joining data from disparate sources to their data warehouse was time-consuming and labor-intensive.

Starburst Enterprise functions as an abstraction layer, allowing doxo analysts to quickly and easily query data in multiple database systems and its data warehouse simultaneously, without requiring ETL or other complex operations.

What Starburst Delivered

- <6 hours per monthly close cycle (from 4 weeks)
- Millions of daily transactions simplified
- Single view into 10 databases

Online food ordering leader improves query performance 10-15x

Overview

This online food ordering and delivery company is one of the largest in the U.S. with over 20 million customers. Its platform connects customers with hundreds of thousands of merchants to place online orders, have their purchases delivered directly to their door, and track orders through a web application.



“Migrating to Snowflake would require two full time engineers and ongoing maintenance, not to mention the time required to stand up the new pipeline. With Starburst, we can query the data where it lives and focus on more important projects for our business.”



Anonymous, Data Platform Engineering Manager

Challenge & Solution

With customer order data residing in both Databricks and Snowflake, the analytics team lacked the ability to easily access both data sets simultaneously. Workarounds led to higher TCO and slower time to insight.

The company deployed Starburst as the single point of access to process over 10,000 queries and 250TB of distributed data daily. Of the company’s analytics workloads, 95% are served through Starburst. The food delivery platform no longer experiences vendor lock-in, as it’s able to federate across all of its data sources in near real-time.

What Starburst Delivered

- 10-15x query runtime improvement
- Significantly lower TCO
- Faster time to insight with real-time reporting and visualizations

Biopharmaceutical company speeds time-to-insight for faster drug development

Overview

This global biopharmaceutical company focuses on advancing innovative medicines to prevent and treat life-threatening diseases, including HIV, viral hepatitis, and cancer. Searching for ways to accelerate its time to market, Gilead aimed to streamline its approach to analytics with a data mesh design.



“With Starburst, our end users don’t need access to the individual data sources directly, they can simply write one SQL statement that reaches out to multiple sources. It makes our end users very productive and the time to derive insights is significantly faster.”



Anonymous, Head of Cloud, Data, and Analytics

Challenge & Solution

This biopharmaceutical company had its analytical data stored across different database technologies such as SAP Enterprise HANA, AWS Redshift, S3, and Oracle Database. This made it challenging to achieve insights quickly and manage data access.

Starburst enables a federated semantic layer and distributed query engine for the data mesh, allowing users to access data quickly and easily, no matter where it lives.

What Starburst Delivered

- 33% increase in time-to-insight
- 500+ data products
- >100 use cases supported

Food & beverage giant accelerates cloud migration

Overview

This food and beverage giant is one of the largest beverage companies in North America. The company is increasingly leveraging analytics and data science to understand its consumer habits and needs and, along the way, unlock new growth.



“The inexpensive, resilient, highly scalable storage of ADLS, along with the security enhancements and analytics anywhere capabilities of Starburst Enterprise are helping to move the company closer to its goal of becoming the global leader in convenient foods and beverages.”



Anonymous, Data Engineering Manager

Challenge & Solution

While transitioning to Azure Data Lake Storage, the company had a problem: roughly 95% of its data remained siloed within on-premises data warehouses like Teradata. It needed a way to analyze both the new data streaming into ADLS and this siloed data.

To assist with the cloud migration off of its legacy data warehouse, the food and beverage giant turned to Starburst.

What Starburst Delivered

- \$200 million in new subscription revenue
- Faster time-to-Insight – from weeks to minutes
- Lower TCO on Azure’s cloud infrastructure



Complement your data warehouse with Starburst.

Today's organizations need reliable solutions that provide the fastest path from data to decision-making. A unified view of disparate organizational data and a future-proof architecture with Starburst are the pathways to better, quicker insights and vendor flexibility. With advanced technology at their fingertips, data leaders can adapt and evolve with the data landscape.

Join the ranks of Comcast, Grubhub, and other companies that are modernizing their data stack by making the warehouse a component of a broader data lake strategy.

The modern data lake is here.
Visit starburst.io to get started.

