snowflake®

CUSTOMER CASE STUDY

PITNEY BOWES USES SNOWFLAKE TO ENSURE CUSTOMER SHIPPING SATISFACTION

INFORMATION TECHNOLOGY AND SERVICES

pitney bowes

COMPANY Pitney Bowes

LOCATION Stamford, Connecticut

SNOWFLAKE WORKLOADS USED

Pitney Bowes is a global technology company providing commerce solutions that power billions of transactions. Approximately 1.5 million clients around the world rely on the accuracy and precision delivered by Pitney Bowes solutions in the areas of ecommerce fulfillment, shipping and returns, cross-border ecommerce, office mailing and shipping, presort services, and financing.

STORY HIGHLIGHTS:

Per-second pricing
Pitney Bowes pays for only the storage and compute it consumes, tracking usage at a granular level.

Instant elasticity
Snowflake's instant elasticity gives Pitney Bowes the compute flexibility it needs to meet a 10x increase in query volume due to COVID-19.

Predictive analytics
Using predictive analytics, Pitney Bowes creates comprehensive shipping dashboards to accurately estimate when a package will be delivered.

> **"Unexpected peaks, such as a global pandemic, happen in the e-commerce shipping world. Thank goodness we had Snowflake to accommodate that volume."**
>
> **—VISHAL SHAH,**
> Solutions Integration & Deployment Architect, Pitney Bowes

CHALLENGE:

Managing costs with an inefficient, inflexible data infrastructure
As a global provider of commerce solutions, Pitney Bowes ingests data from multiple systems, including shipping and transactional data. Its original data infrastructure was supported by SQL Server, an on-premises relational database. That system could not scale to support increasing data needs, and it prevented data sharing. Data existed in silos, and departments focused solely on their own data sets.

As a result, Pitney Bowes transitioned to a public cloud infrastructure. However, that required the DevOps team to spend a lot of time and effort managing compute resources. Granting data access to different teams involved complex management of privileges, preventing the organization from segregating resources across different teams. Complex queries running on the same cluster resulted in frequent downtime.

During peak seasons, such as holiday shopping seasons, they had to pay up-front fees to their cloud provider for expected peak usage.

Pitney Bowes needed a way to scale compute and storage while providing dedicated compute environments to different departments, so one department's queries did not affect another's.

SOLUTION:

Consolidated data and simplified analytics with Snowflake
Snowflake on AWS enabled Pitney Bowes to consolidate its data infrastructure across business units, creating a single source of truth.

Data needs continue to increase, but with Snowflake, Pitney Bowes now efficiently ingests data from 20 different systems across 25 business units. Data is loaded into Snowflake from S3 through ELT tools such as SnapLogic and DataStage.

10x
Increase in query volume year over year

98%
Decrease in time for financial reporting

Zero
Overhead to create new environments

Business units combine multiple data sources (such as shipping, IoT, and transactional data) to create comprehensive dashboards that accurately inform them of a package's journey to delivery. Pitney Bowes can link information about a parcel to the customer receiving it, including customer type, parcel cost, and whether the customer has open cases in the CRM system. Once this data is linked, subsequent users don't have to re-create the same information. All users across the enterprise now benefit from the linkage.

Pitney Bowes also ingests real-time streaming data into Snowflake using tools such as HVR, Apache Kafka, AWS Lambda, and Amazon Kinesis. This provides a near real-time view of the operational process and where parcels are in their delivery route.

> **"We are collaborating much closer with Engineering, Marketing, and other groups to understand what data is being brought into Snowflake. We're sharing more than we ever have before, and the time-to-market for new solutions is so much faster. Before, we were working in silos and now we're working much more across all the organizations."**
>
> **—DIANE BORGIA,** Data Warehouse Technologist, Pitney Bowes

RESULTS:

Meeting delivery commitments with predictive analysis

Before Snowflake, Pitney Bowes analysts manually queried data from multiple sources, then manually linked the data. Now all critical business data is centralized in Snowflake and that data is made available for internal secure data sharing.

Pitney Bowes created predictive dashboards to ensure satisfied customers with its guaranteed package delivery date program. To get an accurate delivery date, Pitney Bowes runs simulations to estimate how long a parcel would take to go from locations A to B and each touch point in between. This analysis includes predictions that match expected volume to available warehouse staff for specific time periods.

Increasing analytics with fast queries

Snowflake's outstanding query performance encouraged more business units to use its platform.

According to Vishal Shah, Solutions Integration & Deployment Architect at Pitney Bowes, "People have been amazed at the query performance of Snowflake. Users have enjoyed the flexibility to write SQL queries instead of doing ETL operations. That's why we've seen such widespread adoption. I'm continuing to onboard one business unit per month, and I'm confident all of Pitney Bowes will soon be on Snowflake."

During the COVID-19 pandemic, Pitney Bowes has seen a 10x increase in query volume year over year. Snowflake enables virtually unlimited computing power for virtually any number of users, providing the performance and actionable insights needed to grow the business.

Freeing resources to focus on higher priority projects

The near-zero maintenance of Snowflake with high availability has lessened the need for dedicated DBA and DevOPs resources.

In the past, 10 to 12 DBAs were dedicated to monitoring the availability and performance of compute resources, including nightly support. With Snowflake uptime and auto-scaling, many of those resources have been freed up to focus on other projects. DevOps in each business unit that were once concerned with bringing in data through CI/CD pipelines can now instead focus on important business data operations.

In addition, there is almost no overhead required to create new environments. According to Shah, "If we need to create a new environment, we basically turn the knob on Snowflake and it creates a node behind the scenes."

> **"An enormous amount of logistical and operational data needs to be considered when shipping a package. With Snowflake removing our data engineering burdens, team spirits have been lifted, and business units have visibility into all the moving parts they couldn't see before."**
>
> **—VISHAL SHAH,**
> Solutions Integration & Deployment Architect, Pitney Bowes

FUTURE:

Using Snowflake to grow the business

According to Shah, "Snowflake will continue to expand its presence at Pitney Bowes." And with centralized usage of Snowflake, data science can focus on enabling business units to grow the business. Eventually, Pitney Bowes will start bringing near real-time data into Snowflake, and with its centralized operational view, it will achieve even higher performance for a better customer experience.