

FLEXPORT BUILDS A SCALABLE DATA MESH WITH SNOWFLAKE AND DBT

LOGISTICS

flexport.

COMPANY Flexport

LOCATION San Francisco, California

SNOWFLAKE WORKLOADS USED



Freight forwarder Flexport exists to make global trade easy for everyone. More than 10,000 clients and suppliers across 200 countries rely on Flexport's software, logistics infrastructure, and supply chain expertise. To provide end-to-end visibility from PO creation to shipment delivery, Flexport ingests and analyzes large amounts of supply chain data.

STORY HIGHLIGHTS:

Architecture supports production workloads

Scalable data mesh that organizes applications layer into domain-bounded data marts with ability to support production use cases.

Extensive network of connectors and tools

From easy-to-use SQL to modeling with dbt, Snowflake supports different domain teams with different skill sets with the tools they need to get work done.

Near-zero maintenance

Remove regular maintenance tasks and automate pipelines with streams and tasks.

CHALLENGE:

A robust data platform needed for unique global trade industry challenges

Migrating data from Flexport's legacy data warehouse to Snowflake improved query performance, reduced time-consuming administrative work, and freed up technical talent to focus on higher-impact work. While other solutions in the market, such as Google BigQuery, are constricted by memory allocation and scales by volume, Flexport needed a solution that could allocate compute

resources on the fly for production and external use cases where latency mattered. "We found that it was hard to guarantee that latency in the other tools we evaluated, except for Snowflake," Flexport's Head of Growth and Analytics, Abhi Sivasailam, said.

Viewing data as a product instead of a byproduct

Flexport's transition to a service-oriented architecture (SOA) presented an opportunity to reimagine the company's data architecture as a data mesh. "This was an organizational shift to federate the creation of rich analytical data assets and govern a broader process for thinking about data as a product," Sivasailam said.

Determining the best way to empower users through decentralization was a key step in Flexport's data mesh journey. According to Sivasailam, "We wanted to lower the friction for the rest of the enterprise to participate in decentralized data product creation, and for us that meant aggressively standardizing on infrastructure."

“**Snowflake is where our source system data goes to be processed and exposed for analytical and scientific use cases.**”

—ABHI SIVASAILAM, Head of Growth and Analytics, Flexport

SOLUTION:

A platform for building a scalable data mesh

Snowflake's multi-cluster shared data architecture enabled Flexport to build a scalable data mesh that organizes application layer data into domain-bound data marts. Flexport's "gold" consumer mart in Snowflake reconciles data from multiple producer marts and powers the majority of Flexport's BI and data science use cases. "Snowflake is where our source system data goes to be processed and exposed for analytical and scientific use cases," Sivasailam said.

5.5x

Increase in active Snowflake users

2 Days

Finance team has been able to bring month-end close processes down by two days

Connecting dbt to Snowflake streamlined Flexport's data engineering workloads. RBAC and column-level security in Snowflake simplified data governance. Snowflake's clean, easy-to-navigate interface, native SQL support, and interoperability with popular BI tools empowered users to explore data with ease.

Snowflake Data Marketplace made it easier to leverage third-party data sets from vendors the company had already implemented, such as Amplitude, and discover new data sets, such as Panjiva. According to Sivasailam, "In some cases, we didn't even know data sets existed until we found them in Snowflake Data Marketplace."

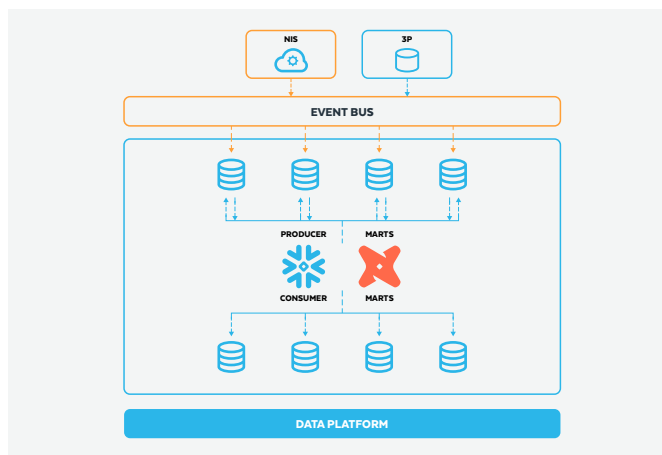


Figure 1. Common data standards

RESULTS:

Delivering data as a product

Successfully implementing Flexport's data mesh strategy with Snowflake has led to a healthier product development process that considers data as a product—not a byproduct. Analytics engineers are essential members of Flexport's development teams and partner with product managers, engineering managers, and software engineers to successfully design, launch, and govern data products.

Decentralization makes it easier for Flexport to identify potential data issues and ensure data quality. "We prefer to drive all of this through data owners, not just one data czar," Sivasailam said. Data governance with Snowflake, through principles of data mesh, provides Flexport greater scalability and confidence to democratize access to trusted data. According to Sivasailam, "Since shifting to data mesh with Snowflake, 5.5x more people are using data across the business on a regular basis."

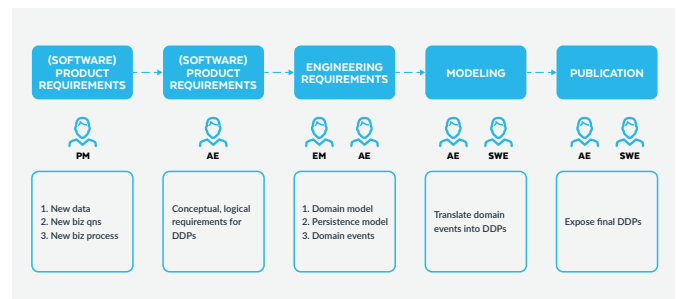


Figure 2. Data product development process

Department-specific consumer marts provide additional flexibility to support users' unique data requirements. For example, Flexport's finance team relies on a dedicated consumer mart that simplifies GAAP revenue recognition and helps bring month-end close processes down by two days.

“When you have a mandate to design data as a product and first-class concern, then your locus of focus shifts to every user.”

—ABHI SIVASAILAM, Head of Growth and Analytics, Flexport

FUTURE:

Sharing data to create new business lines and creating new products

While many freight forwarding competitors are operating in archaic ways, Flexport is data-driven and sits on a wealth of shipping and logistics industry data that has proven out its business's success. Flexport plans to embed its data products and interactive reporting that are developed for internal use into the Flexport Platform. "Snowflake and dbt are on the back end powering what's surfaced to the consumer, mediated by a BI tool," Sivasailam said.

Flexport is also looking to increase its relationship with customers and partners that are interested in the proprietary data and insights it's gathered. Leveraging Snowflake Secure Data Sharing to enable live data sharing with its clients and suppliers is a priority for Flexport. According to Sivasailam, "It's fundamentally about how we get the best data in the world for trade operators and build reliable applications on top of it." The company is exploring platform monetization by providing its data through Snowflake Data Marketplace as well.

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)