

limits of what an on-premises data warehouse could do to meet our business needs. We wanted to move to cloud quickly and open up new analytics capabilities for our teams.

Considering and testing cloud platforms

Teradata had worked well for us when we first deployed it. Back then, Teradata was a market leader in data warehousing, with many of the leading banks invested in it. We chose it for its high-performance analytics capabilities, and our marketing and risk management teams used it heavily. It also worked well with other SAS tools we were using, and SAS remains a good tool for accessing our mainframe.

Ten years into using Teradata, we had a lot of product-specific data stores. It wasn't a fully formed data lake architecture. We also maintain more than 200 SAS models. In 2019, our Teradata appliances were nearing capacity, and we knew they would need a refresh in 2021. We wanted to avoid that refresh, and started doing proof-of-concept cloud testing with both Snowflake and Google Cloud.

When we did those trials, we ran comparative benchmarks for load time, ETL time, performance and query time. Snowflake looked just like Teradata, but in the cloud. With Google, we looked at all the surrounding technology of the platform. We couldn't be on a single cloud platform if we chose Snowflake. We picked Google Cloud, since it would let us simplify and offer us a lot more options to grow over time.

[Latest stories](#)[Products](#)[Topics](#)[About](#)[RSS Feed](#)

with cloud, but ultimately we're aiming not to pay for storage that's just sitting there, not



Cloud won't interfere with other users' performance, but would cost us money. So we're running training to ensure people aren't making those types of mistakes, that they're running the right types of queries.

Shifting to cloud computing

The actual cloud migration involved working closely with the security team to meet their requirements. We also needed to align data formats. For example, we had to make sure our ETL processing could talk to [Google Cloud Storage buckets](#) and BigQuery data sets. We're finding that for the most part the queries do port over seamlessly to [BigQuery](#). We've had to tweak just a handful of data types.

Since moving to cloud, the early results are very promising; we're seeing 3 to 4x faster query performance, and we can easily turn capacity up or down. We have five data marts in testing to use real-world data volumes to get comparison queries.

We're still making modifications to how we set up and configure services in the cloud. That's all part of the change that comes when you're now owning and operating data assets securely in the cloud. We had to make sure that any personally identifiable information (PII) was stored securely and tokenized. We'll also continue to tune cost management over time as we onboard more production data.

Managing change and planning for the future

[Latest stories](#)[Products](#)[Topics](#)[About](#)[RSS Feed](#)



Blog

Menu ▾

focus and executing on it. We're retraining with these new skills, and even finding new roles that are developing. We built a dedicated cloud-native team—really an extension of our DevOps team—focused on setting up infrastructure and using infrastructure as code.

The program we've built is ready for the future. With our people and technology working together, we're well set up for a successful future.



POSTED IN:

[DATA ANALYTICS](#)—[GOOGLE CLOUD PLATFORM](#)—[CUSTOMERS](#)—[FINANCIAL SERVICES](#)—[NEXT](#)

RELATED ARTICLES

◀ Dataflow Under the Hood: understanding Dataflow techniques

Multi-language Dataflow pipelines enabled by new, faster architecture



Follow Us



[Privacy](#) [Terms](#) [About Google](#) [Google Cloud Products](#)

Language ▾



Help