Snagajob Improves Performance 300% with Snowflake

snagajob

CUSTOMER Snagajob

PARTNER Snowflake

snowflake

CASE STUDY

Snagajob is America's largest online portal for hourly job seekers and employers. Its website serves 65-70 million registered job seekers, who seek hourly employment to meet their lifestyle and earnings needs. Snagajob uses business intelligence to connect workers with opportunities based on best-fit skills, their past interests and their previous website visits.

For employers, Snagajob helps narrow down the applicant pool for open positions to a subset of suitable candidates. And Snagajob's delivers additional value by informing employers and job seekers of emerging, real-time and longer term job trend analysis from the website.

Snagajob's real-time analytics inform its sales representatives what interests job seekers – types of benefits, level of compensation, an industry of interest, or even how a company visually presents its employment opportunities. Snagajob also uses an 18-month rolling history of job posting interactions from its database to identify hiring trends so its data analysts and salespersons can identity which industries are out-performing others for advertising and hiring.

A key strategy for Snagajob was to scale the depth and breadth of its analytics while reducing its technology costs.

"I could see the writing on the wall. We couldn't support the growth in our business unless we made a dramatic shift with our IT."

— Robert Fehrmann, Principal Data Architect, Snagajob

THE CHALLENGE

Deploy a Data Warehouse and Analytics Solution that Could Keep Pace

Snagajob's original on-premises data warehouse and analytics solutions emerged from an online transaction processing (OLTP) environment with a monolithic server. The system's significant performance problems meant Snagajob's sales force and its data analysts struggled to get immediate access to data.

They looked at several new approaches and chose a combination of an on-premises, Cloudera distribution of Hadoop and an HP Vertica SQL database for data warehousing and business intelligence. The system had to be able to manage and analyze large amounts of structured and semi-structured data, but it also had to solve the problem of rapidly "getting the data out" to end users. "We thought that we had an answer to our problem," Fehrmann said.

But as Snagajob's business continued to expand, Fehrmann and his staff again found themselves in a familiar place: projected increases in demand would max out this new system in two years. "We needed a system that could grow with us," Fehrmann said. "Not one we'd have to continually reinvent."

Snagajob then considered the potential scalability gains by moving to the cloud. "We were still seeing problems in scalability, cost, maintenance and concurrence. It was around that time we took our first look at Snowflake. Boy, that was a different story!"

SNOWFLAKE'S COST OF PERFORMANCE DIFFERENCE

The three criteria that Snagajob had consistently targeted for its analytics were:

1. Democratizating data so that all users got what they need when they need it.
2. A scalable solution for high and low activity loads.
3. Fast and reliable system performance at low cost.

Snagajob saw immediate improvement in query performance and with high concurrency loads over its existing platform. "With Snowflake, our sales people and our data analysts didn't step on each other anymore," Fehrmann said. "This improved performance significantly."

Improving cost of performance and scalability

Snagajob also needed a cost-efficient and scalable system that could expand and contract and only pay for what they used. "We would see this huge spike at 9am – every single day except for Saturday and Sunday," Fehrmann said. "By noon, usage would drop by half. By 5pm, we hardly had any load on the system."

Snagajob used Snowflake to stage a scalability and cost of performance test. "We ran a full end-to-end system test that covered all peak and valley demand scenarios that our system experienced," Fehrmann said. "We threw a lot of cloud computing at Snowflake at peak times and reduced that infrastructure in times of low demand. What we wanted from Snowflake was better performance for the same money. What we achieved was three times the performance we had before for the same money!"

WHY SNOWFLAKE

Implementing Snowflake's data warehouse in the cloud delivered several important business benefits to Snagajob:

Improvement in query performance and in the handling of high concurrency loads

Snagajob eliminated resource contention between its sales force and data analysts. Their queries run much faster and Snowflake easily met Snagajob's concurrency requirements.

System scalability

Using Snowflake, it was easy for Snagajob to scale data warehouse resources up to meet peak demand periods, and down for periods of lower use later in the day.

Cost of performance

A major objective for Snagajob was getting more performance for the money out of its analytics. By taking advantage of Snowflake's infinite elastic scalability, Snagajob improved cost of performance by 300%.

ABOUT SNOWFLAKE

Snowflake is the only data warehouse built for the cloud. Snowflake delivers the performance, concurrency and simplicity needed to store and analyze all of an organization's data in one location. Snowflake's technology combines the power of data warehousing, the flexibility of big data platforms and the elasticity of the cloud at a fraction of the cost of traditional solutions. Snowflake: Your data, no limits. Find out more at **snowflake.net**.