SOLUTION BRIEF

Fast Track to Analytics: The Cloud-Native Data Warehouse

DATA USERS NEED DATA

To make use of data for reporting or analytics, the data first needs to get into the hands of data users. That's obvious, but it's not necessarily simple. Getting data to users can be challenging with today's legacy data warehouses and big data platforms. Data users face a number of challenges in trying to access the data they need, including:

* ***Difficulty getting access to data:*** Data users spend most of their time waiting for access to the data they need and only a small share of their time using that data to generate business value. This is because there are many disparate sources of data today, and data may be siloed between different departments or organizations.
* ***Complexity of today's data:*** Due to the complex nature of the big data being consumed today (semi-structured data such as JSON, Avro, XML), significant programming skills might be required to extract the data from sources.
* ***Competition for resources:*** Data users frequently end up competing for technology resources and performance with other users in the organization.
* ***Slow time to insight:*** The complexity and inefficiency of current solutions mean that it can take days or weeks for data to be processed, let alone to be used to deliver business insight.

As a result of these challenges, valuable analysis takes too long, if it can be done at all, which leads to missed insights.

> *"We needed a system to support our data scientists and analysts that could keep up with our rapid growth. Snowflake is faster, more flexible, and more scalable than the alternatives on the market."*
> — Craig Lancaster, CTO
>
> JANA

Problem Lies With the Current Options

Ideally, every organization wants the shortest and simplest pipeline from where the data resides to its analysts or other data users. Unfortunately, with current options for traditional data warehouses and big data platforms, this pipeline is long, convoluted, inflexible, and expensive.

* ***Cost:*** The significant time, money, and effort to deploy traditional data warehouses and big data systems lead to limited IT resources becoming the bottleneck to data analysis. Additionally, the upfront costs to purchase and deploy traditional data warehouses are a huge barrier.
* ***Complexity:*** Data warehouses and big data systems are significantly complex, forcing analysts to rely on programming resources outside their control, or learn complex programming languages.
* ***Inelastic:*** Most legacy data warehousing solutions and big data systems are cumbersome to scale up and down. Traditional systems also were designed around the assumption of a fixed amount of available resources that need to be shared among all users and workloads. As a result, as demand on the systems rises, the processing slows.
* ***Limited ability to process diverse data:*** Traditional systems were designed for structured data; big data platforms were designed for semi-structured data. Most solutions require the purchase of both types

Getting data to analysts can be challenging

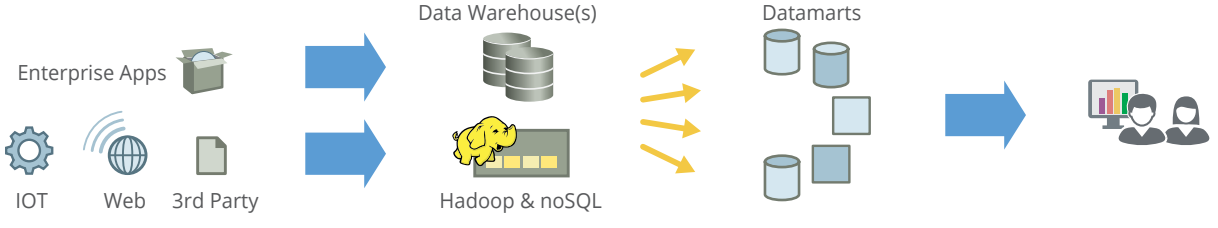

of systems to accommodate both types of data. It becomes the organizations' responsibility to shuffle data among the various types of systems and ensure consistency of data.

REINVENTING THE DATA WAREHOUSE FOR MODERN ANALYTICS

The Snowflake Elastic Data Warehouse takes a new approach to getting data to the data analysts who need it. Built from the cloud up, with no infrastructure to deploy and manage, the Snowflake data warehouse as a service can work with both structured and semi-structured sources. As a result, users can work with all the enterprise's business data and workloads, using standard SQL instead of complex programming languages.

In particular, the Snowflake Elastic Data Warehouse addresses the current challenges faced by data users with the following capabilities and innovations:

* ***Ease of use:*** Snowflake is built to use standard SQL, so it does not require that data users learn new or specialized tools and skills to gain quick, easy access to the data they need.
* ***Ability to accommodate structured and semi-structured data:*** Snowflake can natively store and process diverse data, both structured (i.e., relational) and semi-structured (e.g., JSON, Avro, XML)—all in a single system, without transformation or fixed-schema requirements, and without sacrificing performance, especially for complex workloads.
* ***Zero Management Service:*** Snowflake takes care of overhead and maintenance that would have been the customer's responsibility. The Snowflake architecture separates compute from storage, which allows near-infinite scaling of data, compute, and user concurrency.
* ***Ecosystem:*** Snowflake's partnerships provide native integration with major third-party business intelligence tools; statistical tools like R, ODBC, and JDBC connectivity; and programming languages like Python.
* ***Reduce time to Analytics:*** Users get fast access to data without getting bogged down handling complex programming for Hadoop and similar platforms. By bringing together multiple data types into one system, then making that data accessible using widely familiar SQL, Snowflake enables data scientists, business intelligence analysts, and other data users to stop spending time assembling data and get straight to analyzing and using the data.

SNOWFLAKE BENEFITS BEGIN IMMEDIATELY

Organizations that have implemented the Elastic Data Warehouse have seen 15x–20x performance improvements, and their data users can perform 100 times more queries—all while reducing their costs significantly. Thus these organizations create business value for themselves by simplifying their data pipeline, removing complexity, and making data easily accessible. This enables the analysts to ask more questions and get quicker answers.

Snowflake advantages — cover all your data, all your users and workloads

CAPABILITIES	LEGACY DATA WAREHOUSES	BIG DATA SYSTEMS	SNOWFLAKE
Ease of use for analytics professional	✔	✘	✔
Ability to accommodate structured (i.e., relational) data	✔	✘	✔
Ability to accommodate semi-structured data (e.g., JSON, Avro, XML)	✘	✔	✔
Ability to handle complex workloads without compromising performance	✘	✘	✔
Zero Management Service	✘	✘	✔
Ecosystem partnerships	✔	✘	✔

Rev SB_FTTA_1_x_032016

snowflake

snowflake.net | @SnowflakeDB