

Zero to Insight with the Snowflake Elastic Data Warehouse™

Getting started with tomorrow's data warehouse today



Today's data, and how that data is used, have changed dramatically in the past few years. Data now comes from everywhere—not just enterprise applications, but also websites, log files, social media, sensors, web services, and more. Organizations want to make that data available to all of their analysts as quickly as possible, not limit access to only a few highly-skilled data scientists.

However, these efforts are quickly frustrated by the limitations of current data warehouse technologies. These systems simply were not built to handle the diversity of today's data and analytics. They are based on decades-old architectures designed for a different world, a world where data was limited, users of data were few, and all processing was done in on-premises datacenters.

THE PAINS OF TRYING SOMETHING NEW

Why do organizations still rely so heavily on legacy systems for data warehousing? Because trying an alternative has historically been too painful. Just getting an environment for testing meant installing servers or an appliance in the datacenter followed by time spent configuring software, extracting and transforming a data sample for loading and then translating queries to work on the new system. That's just to try something different—actually deploying a new system adds huge upfront costs and even more challenges—capacity planning, ensuring availability, performance characterization and tuning, setting up monitoring, and more. The harsh reality of data warehousing is that conventional solutions are simply too costly, inflexible, and complex for today's—not to mention tomorrow's—data.

“Pulling insights out of a large volume of data is something that Snowflake has enabled us to do rapidly.”

Tamer Hasan, CTO, WhiteOps

“Big data” platforms do not make things any easier. They require significant work and specialized resources to configure and operate, and they also require learning new skills and tools because they were not built to be data warehouses and to support SQL.

Given the daunting costs, effort, and necessary investments in hardware, software and personnel, it is no surprise that organizations still find themselves largely tied to the limitations of their legacy data warehouses. They find themselves forced to evaluate new projects not based on their potential value to the company, but based on whether they fit their current data warehouse.

SNOWFLAKE REINVENTS THE DATA WAREHOUSE

The harsh reality of data warehousing is that conventional solutions are simply too costly, inflexible, and complex for today's—not to mention tomorrow's—data. These solutions were designed for managing predictable, slow-moving, and easily categorized data that largely came from internal enterprise applications under your control.

They require customers to purchase everything they need for peak demand up front, spending hundreds of thousands of dollars (millions in some cases) just to get started. This all but guarantees that most of the technology

will sit underutilized the majority of the time. As one Director of Analytics put it, “We have to buy for the 99th percentile even though we only reach that level one day per year.

Snowflake developed a completely new data warehouse based on a fundamental vision: enterprises should be able to put all their data in one place, then give access to that data to all their users. To make that possible, this data warehouse needed to be trivially easy to deploy, manage itself automatically, and work with organizations’ existing skills, tools and applications.

From their years of experience with current solutions, Snowflake knew that this could not be achieved by starting with existing architectures or code bases. We saw that we had to start over by building a new data warehouse from the ground up, one designed from the start to make it easy for users to focus on getting value from data, not on managing infrastructure.

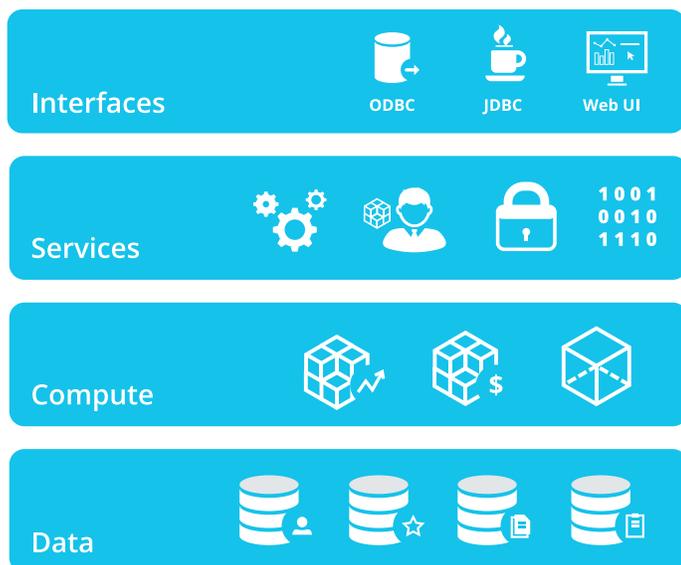


FIG 1 Built from the ground up for the cloud, Snowflake’s unique architecture physically separates and logically integrates compute and storage

MAKING CHANGE SIMPLE AND PAINLESS

The result: a data warehouse that is both easy to try and easy to deploy. You can simply load data and start querying—no hardware to get ready, no software to install and configure, and no tuning required. Simply “load and go” to start using your data.

Snowflake’s patent-pending architecture and technology eliminate the key obstacles to making use of data:

- * **Eliminate infrastructure** Snowflake is delivered in the cloud, so there’s no hardware or software to procure, install, or configure.
- * **Accept your data the way it is** Because Snowflake natively stores and processes both structured and semi-structured data (e.g. JSON, Avro, and more), you do not need a complex data pipeline to make your data ready for Snowflake—load it in its native form, even without defining a fixed schema for your semi-structured data.
- * **Leverage existing SQL skills and tools** Snowflake is built for standard SQL. That means that your analysts don’t need to learn new skills, and the broad universe of tools that speak SQL—the same tools you’re using today—work with Snowflake.
- * **Let the data warehouse handle the gruntwork** Snowflake is designed as a cloud service. That means that it automatically takes care of all of the administrative tasks that other data warehouses put in the way. No manual tracking of what has been loaded, no picking distribution keys, no maintenance of indexes, no vacuuming, no queue management—it is all taken care of by Snowflake.

Scale elastically, at any time: Snowflake’s patent-pending architecture can scale up and down storage, compute, and users at any time, without disruption or downtime.

That architecture allows you to load and query at the same time without performance degradation, and it allows you to store any volume of data at a low cost. You don't need to worry about how to minimize the volume of data in the data warehouse, and you don't need to worry that you'll be unable to meet demand if your capacity planning was not exactly right.

"Snowflake is the first analytic database that really leverages the power of the cloud."

Jeff Shukis, VP Engineering and Tech Ops, VoiceBase

GETTING STARTED WITH THE SNOWFLAKE ELASTIC DATA WAREHOUSE

It's easy to try out Snowflake. What you need to do:

- * **Get your data to the cloud** If you already have data in cloud storage services like Amazon S3, Snowflake can load it directly from there. Otherwise, you can use Snowflake's tools to move the data that you want to load into Snowflake to the cloud.
- * **Load your data** Quickly and painlessly load your data into Snowflake. Snowflake's unique architecture means you can apply whatever resources you need to load data as quickly as you need, without impacting performance of other jobs. Because Snowflake natively understands semi-structured data like JSON

and Avro, you can load that data directly, without transformation and without needing to define a fixed schema for it.

- * **Start analyzing data** You can create and submit standard SQL queries using the built-in SQL editor in the Snowflake web interface. Using Snowflake's ODBC and JDBC drivers, you can also connect a wide variety of SQL and BI tools. Because Snowflake automatically optimizes data and queries, you don't have to worry about managing data distribution, indexes, cleanup, or metadata updates.

START GETTING MORE OUT OF YOUR DATA

Breaking free from the shackles of conventional data warehouses has never been easier, less expensive, or less risky. Instead of avoiding new projects because they do not fit your current data warehouse, use Snowflake to easily and inexpensively support any scale of data, compute, and users that those projects require.

With Snowflake, say goodbye to yesterday's challenges of physical infrastructure and complex systems requiring highly-skilled, specially trained staff to use them. Instead, start getting the full value out of your data--in a fraction of the time and at a fraction of the cost of your old system.

TRY IT OUT NOW

To try out the Snowflake Elastic Data Warehouse Service, sign up for a free trial on our website at www.snowflake.net, or contact us at customer@snowflake.net.

About Snowflake

Snowflake Computing, the cloud data warehousing company, has reinvented the data warehouse for the cloud and today's data. The Snowflake Elastic Data Warehouse is built from the cloud up with a patent-pending new architecture that delivers 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. The company is backed by leading investors including Altimeter Capital, Redpoint Ventures, Sutter Hill Ventures and Wing Ventures. Snowflake is headquartered in Silicon Valley and can be found online at snowflake.net.