

## COST REDUCTION AND SIMPLIFICATION



- Single data platform with cloud robustness
- A data platform that's both a data lake and data warehouse
- A simplified data architecture that complements your existing cloud strategy



- Patented micro-partitioning eliminates the need for developers to implement indexes, partitions and statistics
- Single platform for natural ingestion of structured and

### Challenges:

As organisations invest in their data infrastructure, as well as their platforms and tools, the complexity, maintainability and cost control become unmanageable.

As an example, a single organisation may have a mixture of on-premise and cloud storage with HADOOP clusters for unstructured data processing, as well as a SQL data warehouse with curated structured data marts.

### Solution:

As storage, compute and memory have become relatively affordable in the cloud, its infrastructure and services have become a viable solution to organisations of all sizes. Featuring the flexibility of being able to scale compute and storage on demand within minutes, with the reassurance of most major cloud vendors guaranteeing 99% uptime on their infrastructure.

This has allowed the data community to also take advantage of cloud computing on major commercial providers with data frameworks and query engines that have adopted or been built on cloud infrastructure from the ground up.

This in return, provides organisations with ready-to-use infrastructure, application development, hosting and storage as a service.

### How Snowflake can help?

semi-structured  
data.  
A hybrid data  
warehouse and  
data lake

**Snowflake**, a cloud data platform, is offered as a software-as-a-service (SaaS). A simple, single interface with a common language in SQL allows you to perform tasks, administration, data definition, data movement, backup and data recovery, as well as advanced capabilities such as data sharing and looking at historical data with Snowflake's proprietary 'Time Travel' feature.

## Simplicity



*“That we don’t need to do  
any configuration or tuning is  
great--we can focus on analyzing  
data instead of on managing and  
tuning  
a data warehouse.”*

*Craig Lancaster, CTO*

**J A N A**

As a 'software as a service' patching, data management, uptime and availability are standard features, even as part of Snowflake's base-level standard edition.

Snowflake tables do not indexes, statistics and partitions when a data engineer builds pipelines for target data artefacts. Unlike its competitors, Snowflake's patented micro-partitioning ensures tables are performant automatically, where a data modeler and engineer will not have to cater, just load data and it's ready for consumption

Ingestion of structured and semi-structured data follows the same ingestion pattern and is just as performant regardless of whether your source data is from on-premise



databases, logs or IOT devices.

Snowflake natively ingests traditional structured data such as CSV, as well as semi-structured formats such as JSON, XML, AVRO and Parquet.