

acceldata

The Snowflake Data Experience

A Survey of Snowflake Users and
How They Optimize Their Data

MARCH 2023



Introduction

Snowflake is one of the most popular cloud data warehouses today. In just one decade, the company has grown to more than 6,800 enterprise customers and \$1.2 billion in annual revenue, more than double its prior year.

Like other cloud data warehouses, including Databricks, Amazon RedShift, and Google BigQuery, Snowflake boasts an attractive combination of low start-up costs, constant innovation, and “it just works” manageability. And judging by its breakneck growth and enthusiastic customers, Snowflake delivers on these features better than its rivals, especially because of its high availability (with near-zero administration), and instantly-deployable infrastructure.

As the Snowflake footprint grows, it's helpful to understand how it is being used and the impact it's having.

Censuswide recently conducted a survey of more than 200 data leaders (which includes Chief Data Officers, VPs of Data Platforms, Data Engineers, and a variety of other titles, all of which are provided below) from across the United States. Perhaps not surprisingly, these leaders, all of whom manage sophisticated, complex data stacks, report serious concern about lack of visibility. The results highlight specific, common concerns, as well as insights into their data investment plans.

Here's a snapshot of our survey respondents, their top concerns, and insight into their Snowflake environment:

Cost is the biggest concern. Most have been using Snowflake for 1-3 years, most have annual contracts, and contract size is typically between \$250,000 - \$1 million. The Snowflake environment for most is between 1 - 50 warehouses. Data teams of respondents typically have fewer than 50 data engineers. The majority of respondents use other data platforms.

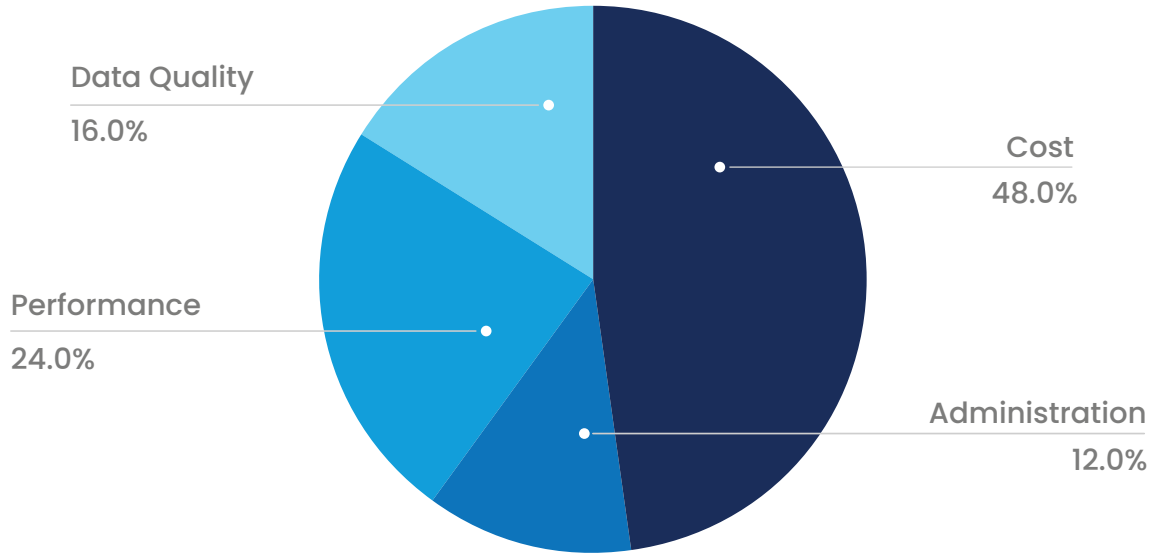
Let's take a closer look at the survey responses:



Top Concerns

Clearly, cost is top of mind for most Snowflake data teams. What's notable about this particular metric is that other top concerns – data quality and performance – are both intrinsically related to cost. When an organization has consistent data quality, they usually see a dramatic improvement in performance, and those things can support better resource efficiency and an improvement in cost to value.

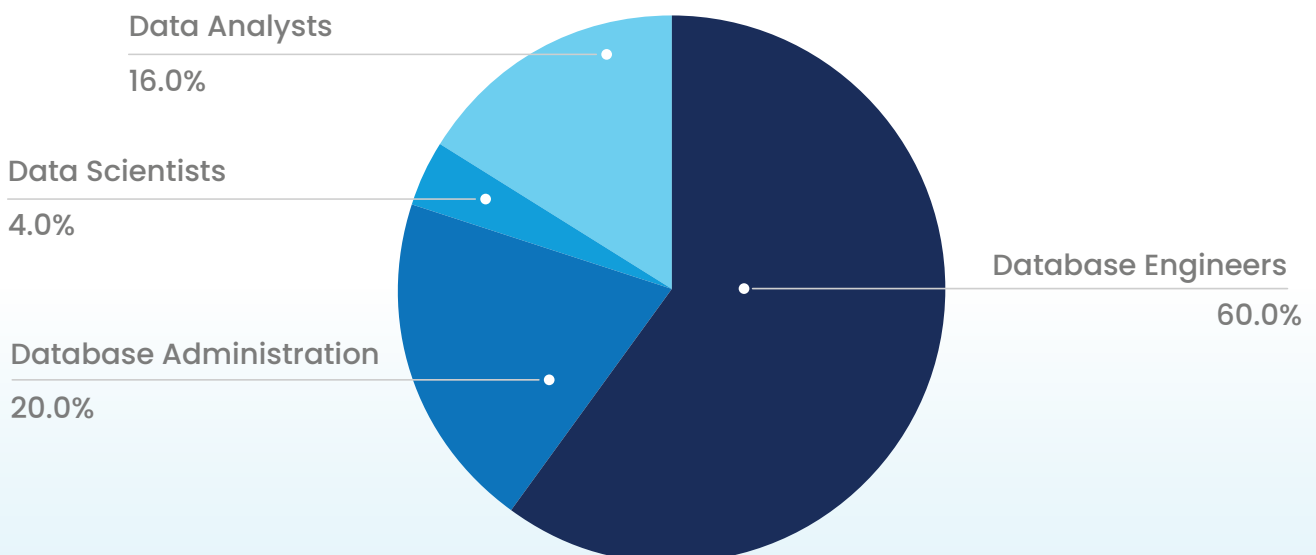
What is your biggest concern?



Snowflake User Personas

It takes a variety of roles to manage Snowflake environments, but it's clear that data engineers make up the vast majority of Snowflake users.

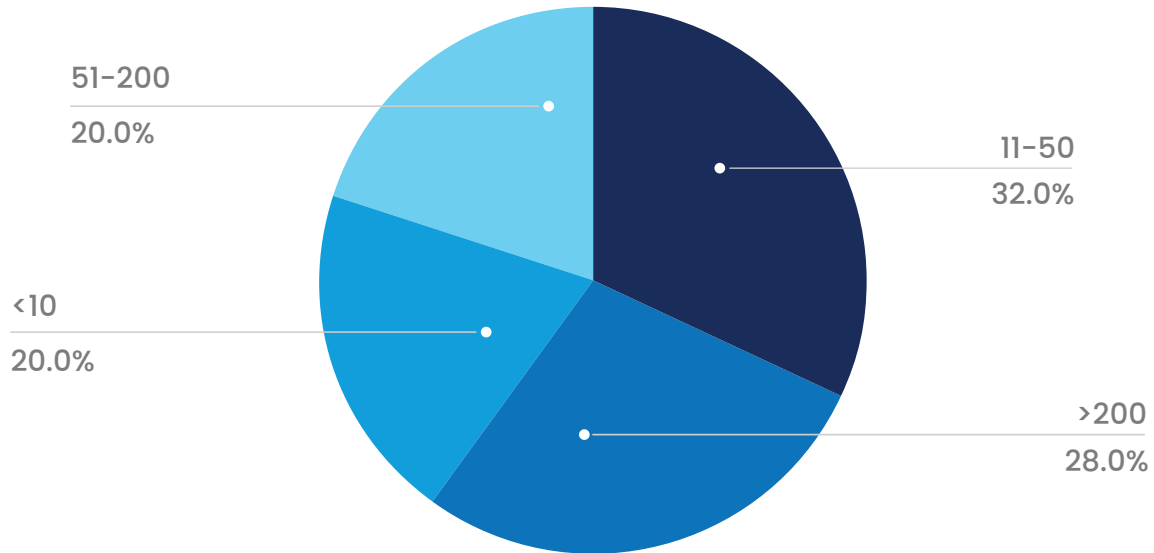
Who interacts with your Snowflake instances?



Number of Snowflake Users

Here again, we have an indication that as organizations grow, their Snowflake environment grows with them. As such, we see that almost half of respondents have more than 51 Snowflake users in their organization.

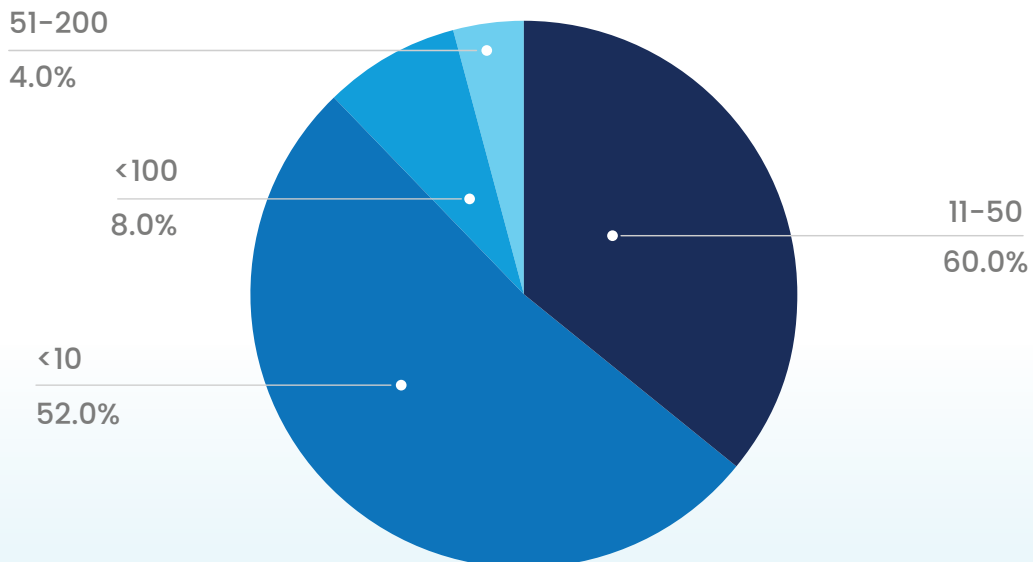
How many Snowflake users do you have in your company?



Active Warehouses

The numbers here show that about half of respondents have fewer than 10 warehouses, while half have more than 10. About 36% have between 11-50 warehouses.

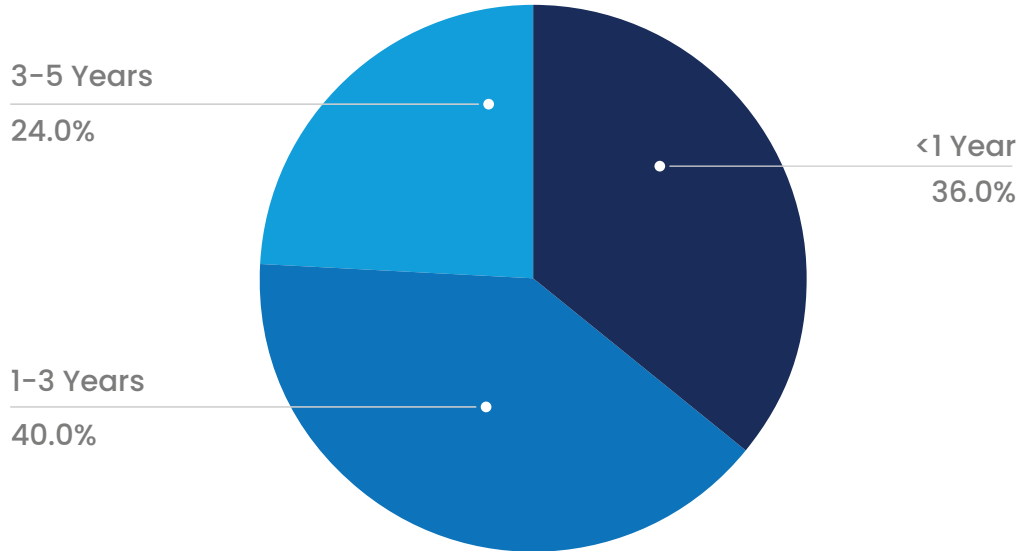
How many active warehouses do you have?



Length of Time Using Snowflake

All respondents have been users for five or fewer years, while 36% have less than one year of using Snowflake at their company.

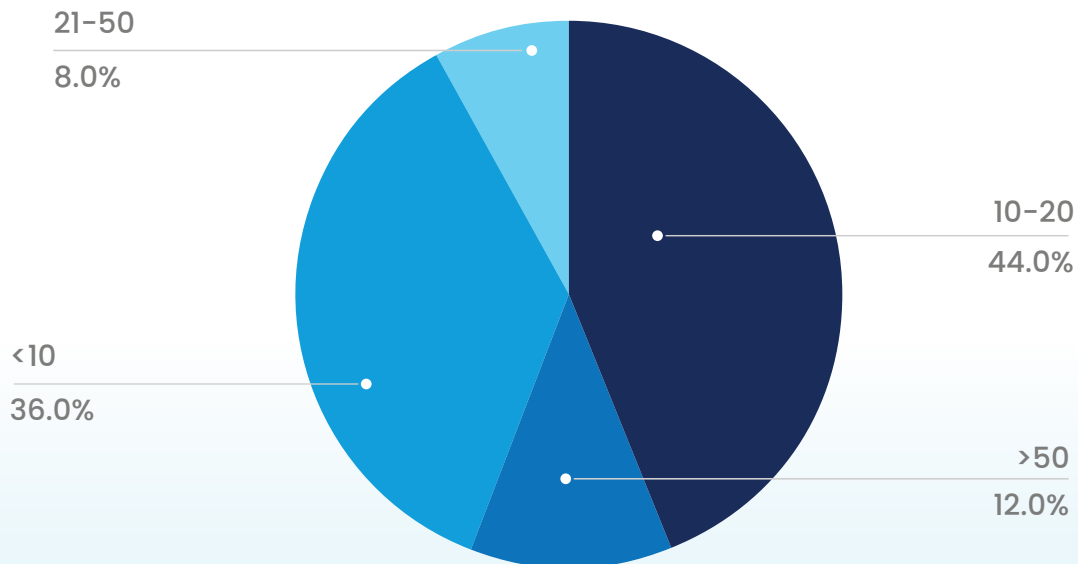
How long have you been using Snowflake at your company?



Size of Data Team

There was considerable variation among respondents in terms of team size, with about 20% having teams with more than 20 members.

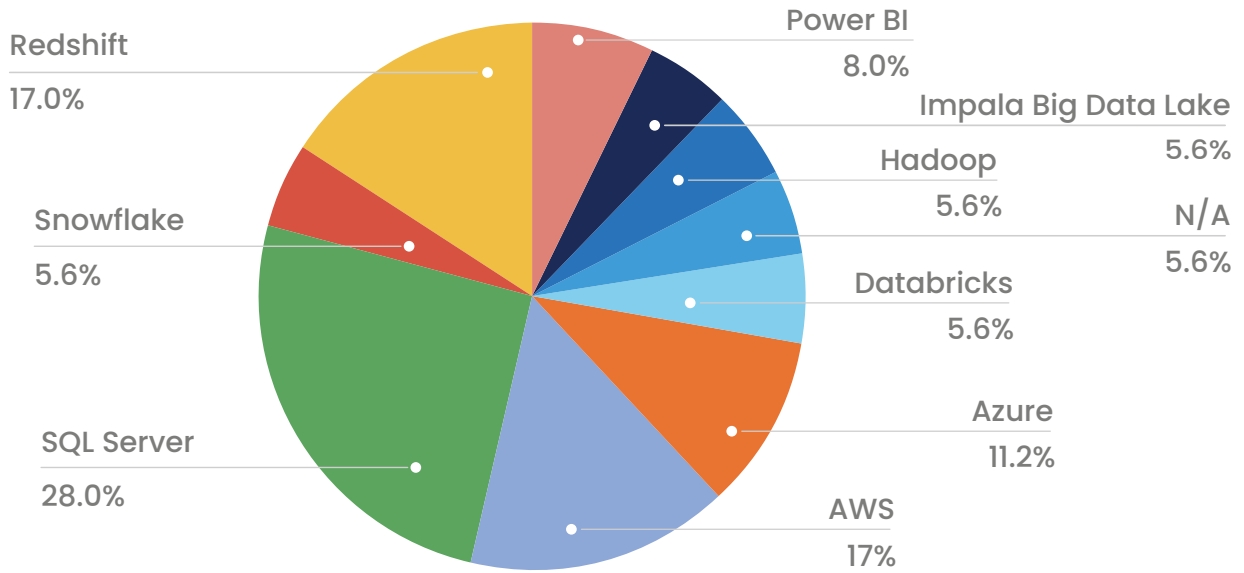
How big is your engineering team?



Data Platforms

A large number of data platforms are used by respondents, as this graph shows.

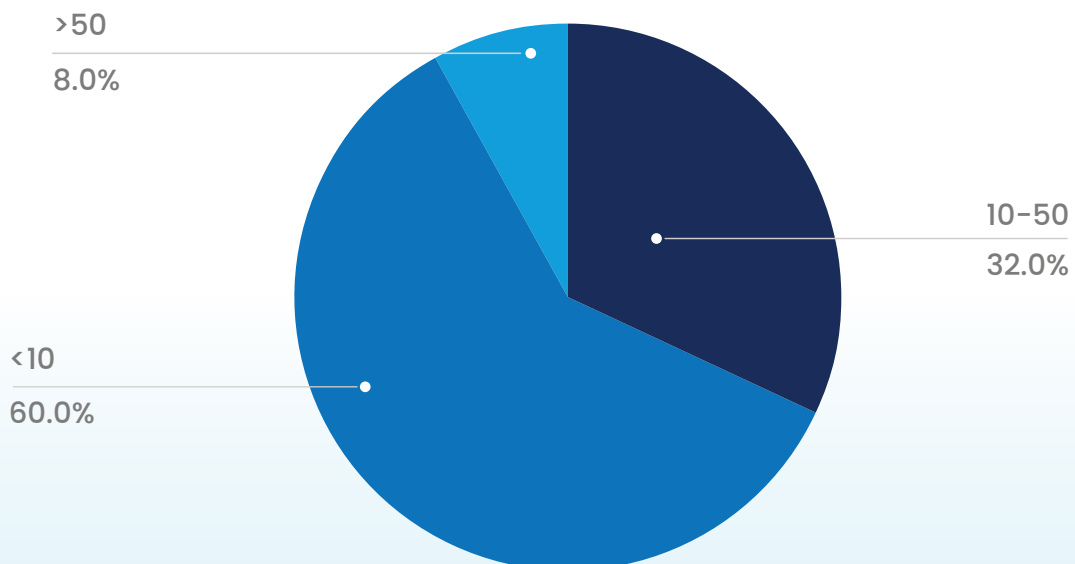
How long have you been using Snowflake at your company?



How Many Applications Use Snowflake?

More than half of respondents use Snowflake for less than 10 applications. We are seeing a trend in increase of applications

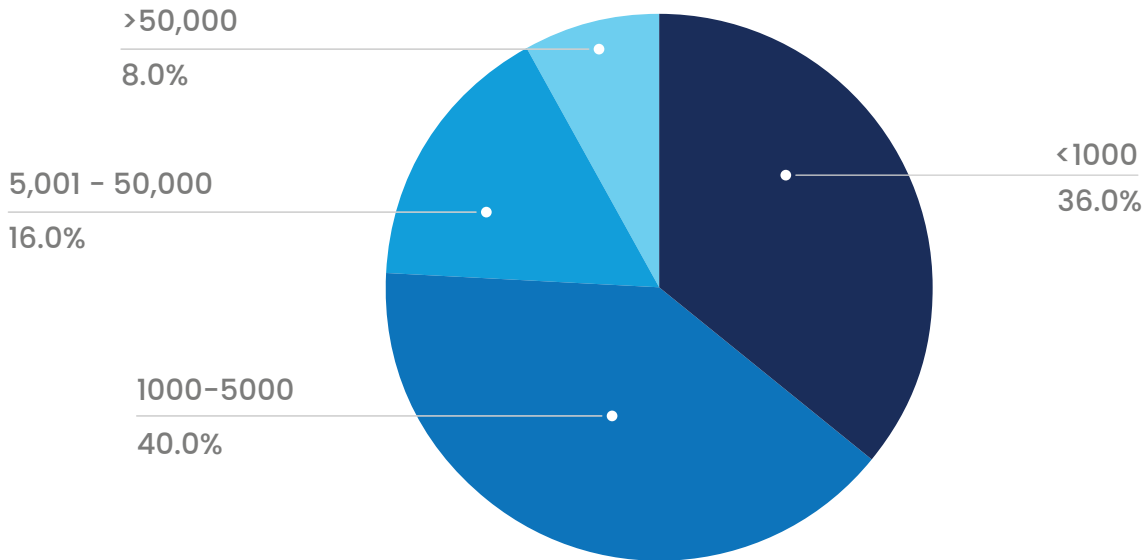
How many applications in your environment use Snowflake?



Daily Snowflake Queries

About a quarter of respondents perform more than 5,000 queries daily. Within that group, 8% run more than 50,000 daily queries.

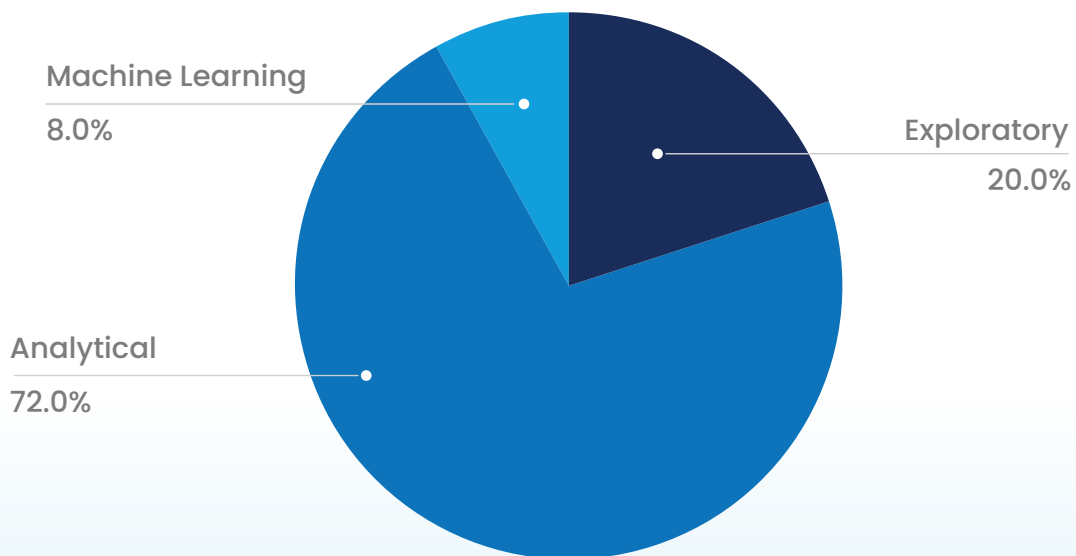
How many queries do you run on Snowflake every day?



Types of Workloads

The survey indicates a wide range of workload types being run on Snowflake.

What workloads do you work on Snowflake?

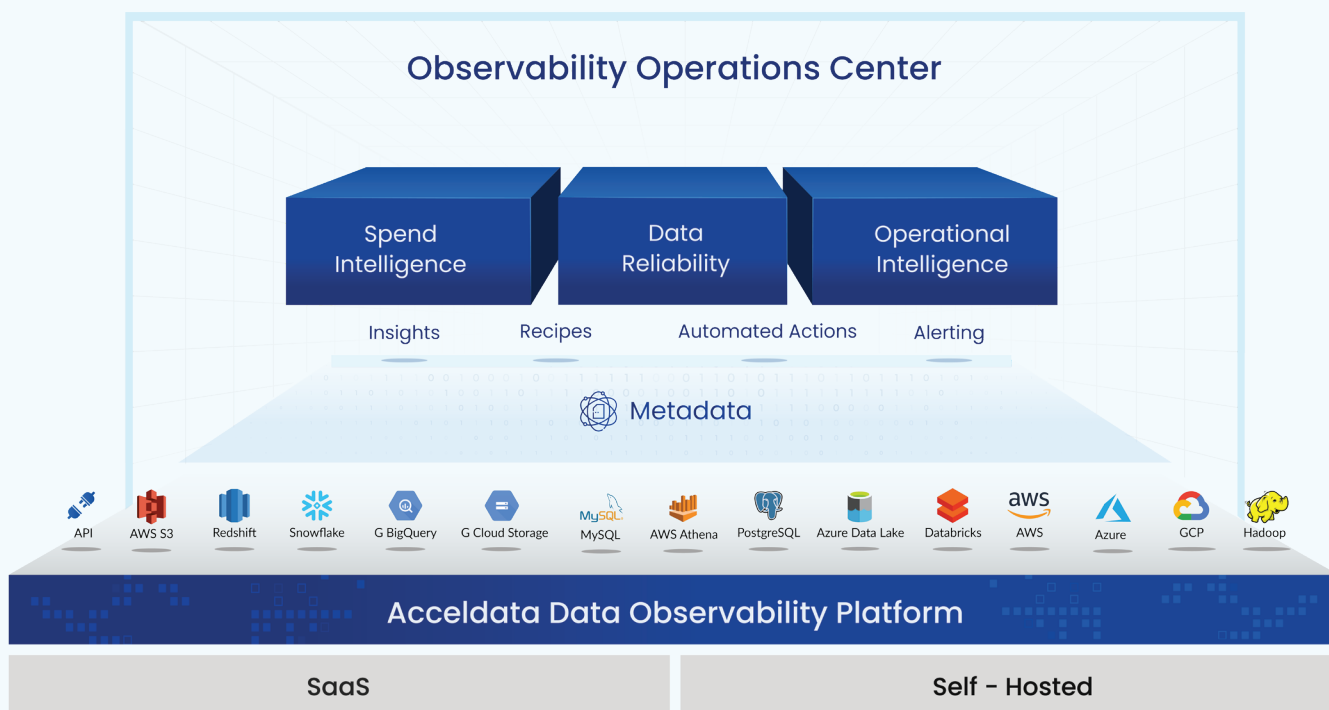


Acceldata: Data Observability For Your Snowflake Environments

The great thing about the cloud is elasticity - you can pay for what you need and optimize your cloud resources to fit your budget and computing processes. The bad news is that your cloud costs can quickly spiral out of control if you don't have real-time insights to effectively manage your resources.

For data to be valuable to a business, it needs to be reliable, accessible, and trustworthy. Data observability provides the real-time insights into your data assets, pipelines, and systems to effectively manage your Snowflake resources and costs.

With Acceldata, you get:



Spend intelligence - insights into resource utilization and associated costs to manage contract plans, current and projected spend analysis, department-level tracking, budgeting, and chargebacks

Data reliability - a full suite of automated data quality and data pipeline monitoring tools to ensure that your data health is at the highest level, alerts when incidents occur, and multi-layer data for root cause analysis.

Operational intelligence - AI-driven insights that show data cadence and resource utilization, offer guardrails for Snowflake best practices, and provide recommendations on how to optimize your environment.

Acceldata provides the visibility and insight to raise your data's efficiency and reliability and help manage your costs. With Acceldata you avoid the hangover effects that often accompany businesses when they switch to cloud platforms.

Learn about the Acceldata Data Observability platform and how it can help you maximize the return on your Snowflake investment with insights into performance, quality, cost, and much more.

[Request A Demo](#)

About Acceldata

Acceldata is the market leader in enterprise data observability for the modern data stack. Founded in 2018, Campbell, CA-based Acceldata, enables data teams to build and operate great data products, eliminate complexity, and deliver reliable data efficiently.

Our Clients



ORACLE



true

dun & bradstreet

& many more

acceldata

© 2023 Acceldata, Inc.

www.acceldata.io