



**PROOF-OF
CONCEPT GUIDE**

TEST-DRIVING SNOWFLAKE:

**THE DEFINITIVE GUIDE TO
MAXIMIZING YOUR FREE TRIAL**



**CHAMPION
GUIDES**

USE YOUR FREE CREDITS WISELY

Snowflake can help your organization get the most out of its data. This guide offers tips and advice, and it's aligned with key test milestones to maximize your 200 free credits, a \$400 value.

Since Snowflake provides usage-based, per-second pricing, you already have the most efficient platform available. But you'll still want to use your credits effectively. It's possible to use all of your free credits in a day simply by leaving an idle virtual warehouse (our compute clusters) up and running. Therefore, it's best to spend a few hours a day for a week loading data and executing queries with Snowflake's auto-suspend feature set to 1 minute or less to help conserve credits. With this approach, you'll get a good feel for the system and still have credits in reserve to get started on your next project.

For most business cases, following the advice in this guide will give you a strong indication of how Snowflake can fulfill your business requirements by the time you have consumed all your free trial credits. If you are interested in purchasing a capacity contract or professional services, please **contact Snowflake**.

GETTING STARTED

First things first...

Before you spin up your first virtual warehouse (compute cluster), get a sense of how Snowflake works and what your data priorities are. Plan how you will use your credits wisely to ensure you get a good look at Snowflake’s strengths and how it can transform your organization. If possible, consider planning test scenarios using a subset of your data. The three most common places to start are:

- 1 **Load your data:** Bring your own data into Snowflake from CSVs, cloud providers, and more. Learn more about loading options [here](#).

Rule of thumb: No more than 25% of your free credits should be spent on data loading. If you think data loading may consume more than 50% of your credit allocation, please contact Snowflake for assistance.
- 2 **Start querying:** Explore ready-to-query data sets from Snowflake Marketplace—see free ones [here](#).
- 3 **Run benchmarks:** Use TPCB benchmarking to see Snowflake’s performance with sample queries and pre-loaded sample data sets. To run benchmarks, simply open a new worksheet and open Tutorials. You’ll see four pre-built options.

Want to tailor your plan to your own needs and get a strong understanding of Snowflake basics? Here are five best practices to follow before you fire up Snowflake:

- 1 Read about Snowflake architecture [here](#).
- 2 Think about which of Snowflake’s capabilities most closely map to the needs of your business case and make sure you test against those needs. Here are some examples: concurrency, ease of use, elasticity, usage-based pricing, SQL-based query capability, flexibility, speed and compatibility with existing ETL (extract, transform, load) and BI (business intelligence) tools.
- 3 Also, think about how you want to divide workloads to cover all the functions you want to test, because loading data, querying, and running external tools (ETL, BI) all consume credits.
- 4 To start using SnowSQL, the Snowflake command-line client, follow the “Snowflake in 20 Minutes” getting started tutorial [here](#).
- 5 Tour the new web interface and learn how to quickly query, visualize, and share insights [here](#).



QUICK TIP

Account billing and usage is easy to access within Snowflake. Always check usage during the trial period to monitor your credits.

TIPS TO USE YOUR CREDITS WISELY

- Pick a reasonable batch of data to load for testing—enough to run a few sophisticated queries and test ETL (ingestion), BI, and concurrency, but not enough to consume too many credits. If you are trying to run production-type workloads, this will generally exceed your free credits.
- As a baseline, try to keep loaded data **under 150GB**.
- **Do not** turn off auto-suspend. By default, the auto-suspend threshold is set to 10 minutes. You should consider **reducing the auto-suspend threshold to 1 minute** during the trial period, especially if you are running regular ETL jobs.
- Account metering and billing runs in real time as you size clusters up and down and run queries. Users with the ACCOUNTADMIN role can view monthly and daily credit use for all warehouses in your account in the Usage tab or with SQL. To view credit use for your account, click Admin > Usage.

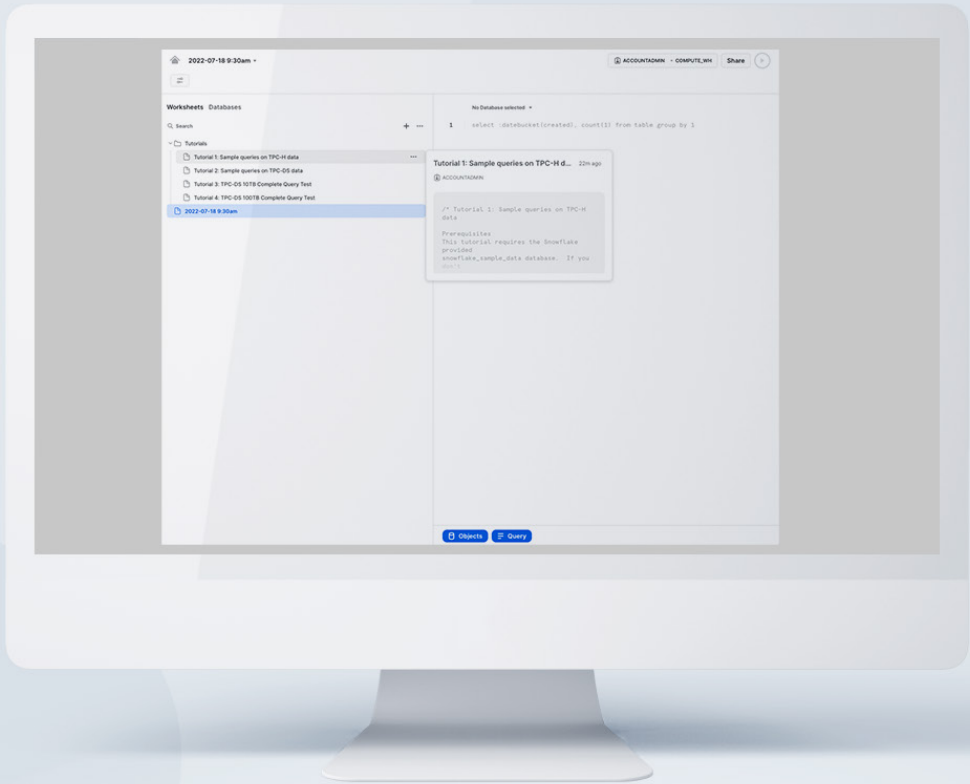
LOADING DATA

REQUIRES 50-100 CREDITS DEPENDING ON THE DATA SET SIZE

STEP 1: CREATE A DATABASE AND TABLE

First, create a database and table in a Worksheet.. The columns and data types in your table need to match the text file you are loading. [Click here](#) for more detailed instructions.

STEP 1
Easily create tables by navigating to a database and clicking into a schema to create tables, views, stages, and more.



QUICK TIP FOR DATA LOADING

For an example with sample code, walk through the loading of Citibike city in our Quickstart Guide [here](#).

STEP 2: FIRE UP ONE OR MORE WAREHOUSES

To load data, you need to create or select a virtual warehouse. Here's how to spend your credits wisely when loading and eventually querying data:

- **Start small.** With Snowflake, you can activate one virtual warehouse or a cluster of them depending on data size, query speed and query or user concurrency needs. You can move your cluster size up and down, and in real time, based on your needs at any given moment. Start with an Extra Small (XS) size (one credit), which will fire up instantly.
- **Grow as needed.** Add or subtract t-shirt sizes (roughly equivalent to cluster nodes) as needed (each active node consumes one credit per hour, so be judicious). If possible, keep the virtual warehouse size for the proof-of-concept period to around eight nodes or less and save Large-size testing for later on. For more information on virtual warehouse sizing, [click here](#).

QUICK TIP

Pick a reasonable batch of data to load for testing—enough to run a few sophisticated queries and test ETL, BI and concurrency, but not enough to consume too many credits. We recommend less than 150GB.

STEP 3: PREPARE TO LOAD DATA

You can load data into Snowflake in several different ways:

- **Use Partner Connect.** Partner Connect simplifies data loading through pre-built integrations with Snowflake's technology partners such as Fivetran, Aloomo and Stitch. Using Partner Connect, you can avoid the manual work associated with creating databases, creating warehouses and separately developing your data pipeline. [Here's a video to learn more.](#)
- **Use Snowsight.** To manually load small files and even spreadsheets from your desktop, the easiest method is to navigate through the Database tab in the interface. Learn more about managing databases and objects [here](#).
- **Manually load bulk data.** You can also load bulk data from [Amazon S3](#) by using [SnowSQL](#), a command-line client. You can find more information about loading bulk data [here](#).

QUERYING DATA

REQUIRES APPROXIMATELY 75-100 CREDITS

Once you load your data, you’re ready for the fun part: running queries. Using the Snowflake sample data set or your own loaded data, you can start running test queries to see the power of the platform. Here are six tasks for querying in Snowflake.

STEP 1: GET COMFORTABLE

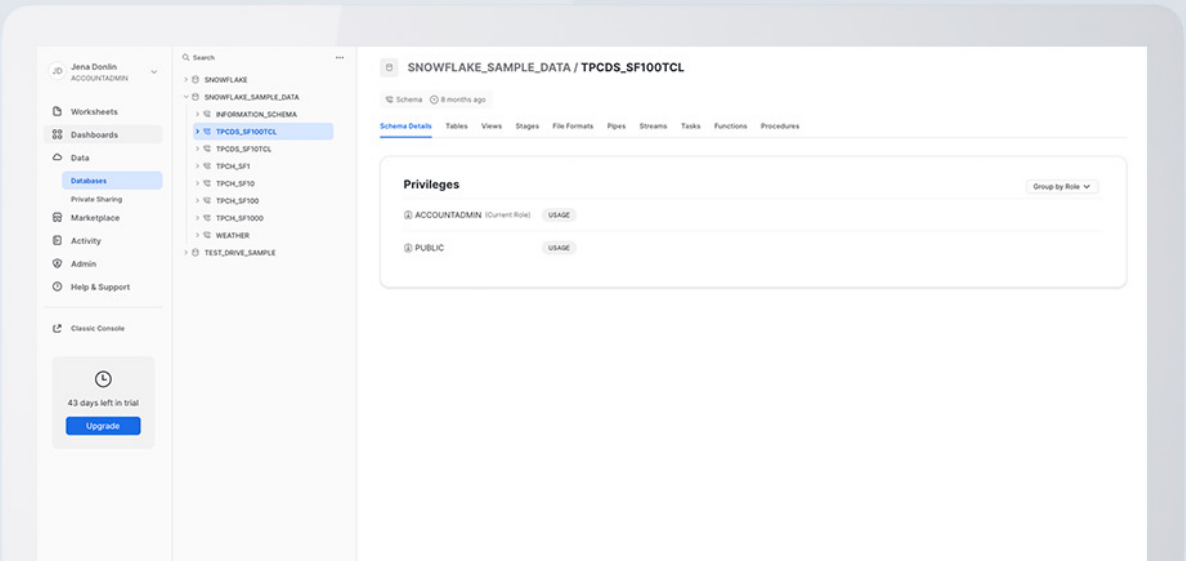
Browse around and look at sample query results. Once you fire up a warehouse, you can also run basic queries on sample data. But keep in mind, though, that sample data queries still consume credits.

The following Snowflake sample data sets and tutorials are available for benchmarking:

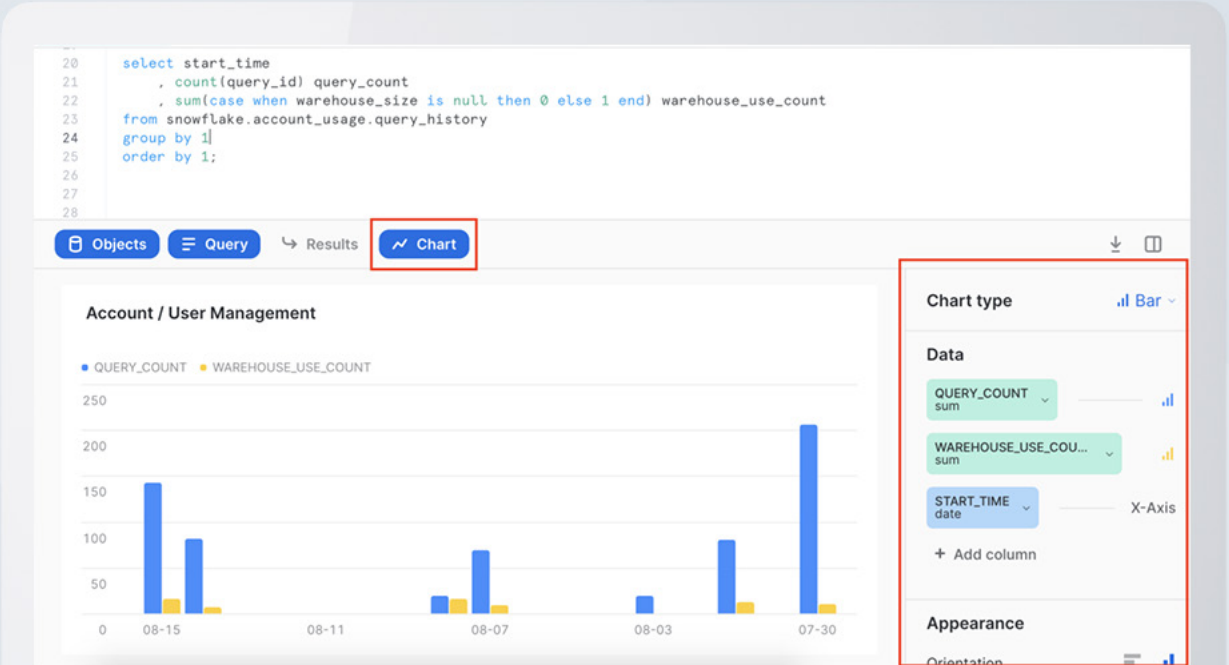
- **TPC-DS** models the decision-support functions of a retail product supplier. The supporting schema contains vital business information, such as customer, order and product data.
- **TPC-H** is a decision-support benchmark. It consists of a suite of business-oriented ad hoc queries and concurrent data modifications.

QUICK TIP

STEP 1
Snowflake sample data sets can be found in the **SNOWFLAKE-SAMPLE_DATA** database.



STEP 2 Sample worksheet.



“Since Snowsight is designed to support data analyst activities, it automatically generates charts from the query result, which will enable us to instantly visualize our data.”
CLARK PERUCHO, Snowflake Data Superhero and DBA

STEP 2: START RUNNING QUERIES ON YOUR OWN DATA INSIDE SNOWFLAKE

Identify workloads and key business-critical or highly common queries you’re currently running and set and run those queries in Snowflake.

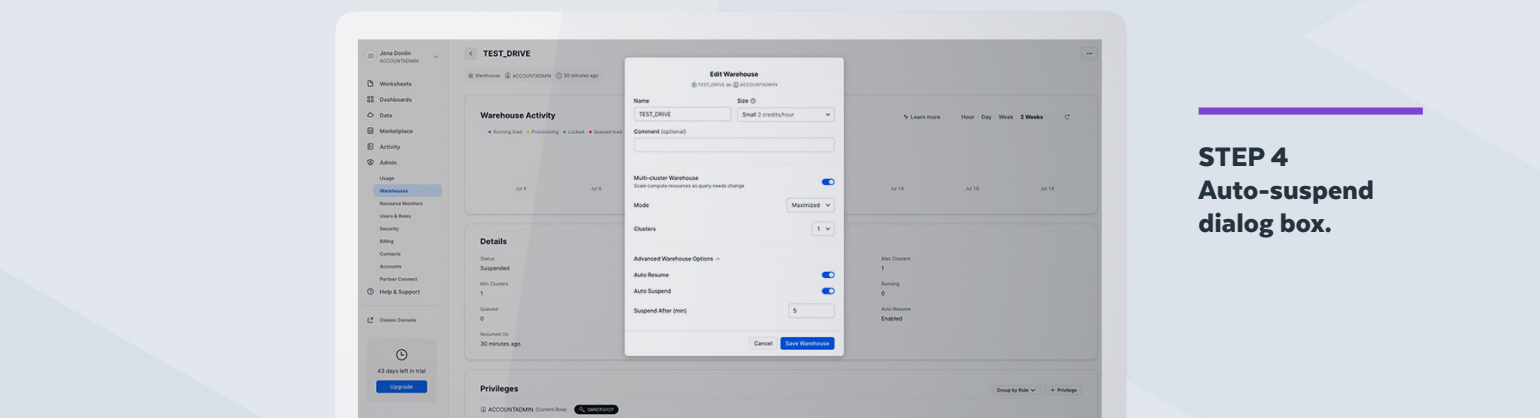
For information on how to run queries, see the following:

- [Overview](#)
- [Using worksheets](#)

STEP 3: TEST QUERIES WITH THE GOAL OF IMPROVING CURRENT PERFORMANCE

Since you are evaluating data platform options, you are likely searching for improvements over your current processes. Make sure you save enough credits to test for the following:

- **Ease of use and speed testing:** How easy is it to set up a query and how fast will it run? Remember, with Snowflake, you can size up virtual warehouses to speed up querying time and you can also choose which warehouse to use to run specific queries. If results are coming in too slow, add some power to the query but, as always, be credit-conscious when you do.
- **Concurrency:** Running multiple queries at once to test query and/or user concurrency results? Similar to the point above, you can increase or reduce the size of your virtual warehouse to change the speed of concurrent queries.
- **Multi-cluster warehouses:** Can you scale compute resources to manage your user and query concurrency needs as they change, such as during peak and off hours? For example, if you are testing ETL feeds or want to emulate time periods where concurrent use is very high, you can set up a multi-cluster warehouse to accommodate these extra loads and then automatically scale back down when peak use subsides. Keep in mind that a multi-cluster configuration will consume more credits. For more information on multi-cluster warehousing, [click here](#).
- **Prototyping:** Can you create and share visualizations? Head over to Dashboards to create and share insights and quickly develop prototypes.
- **Extensibility:** Can data scientists, data engineers, and developers work in the language of their choice? With Snowpark, they can access the secure and governed Snowflake platform and use the elastic performance engine while working in their preferred language.



STEP 4
Auto-suspend
dialog box.

STEP 4: SET UP AND RUN ETL TOOLS INSIDE SNOWFLAKE

Test your ETL processes. If you set up regular interval ETL runs, make sure you change the auto-suspend default from 10 minutes to 1 minute to save credits.

STEP 5: TEST EXTERNAL BI TOOLS

Save some credits for connecting and testing any BI tools you’re currently using (such as Tableau). [Click here](#) for more details.

Put aside some time for connector setup, keeping in mind that every vendor or architecture has slight variations in BI syntax. Convert a couple of queries to BI visualized reports to get a feel for how your BI works inside Snowflake. But remember, it takes credits to configure and run BI reports so don’t run too many reports; you can run a few test reports using a only handful of credits.

QUICK TIP

Convert a couple of queries to BI visualized reports to get a feel for how your BI works inside Snowflake.

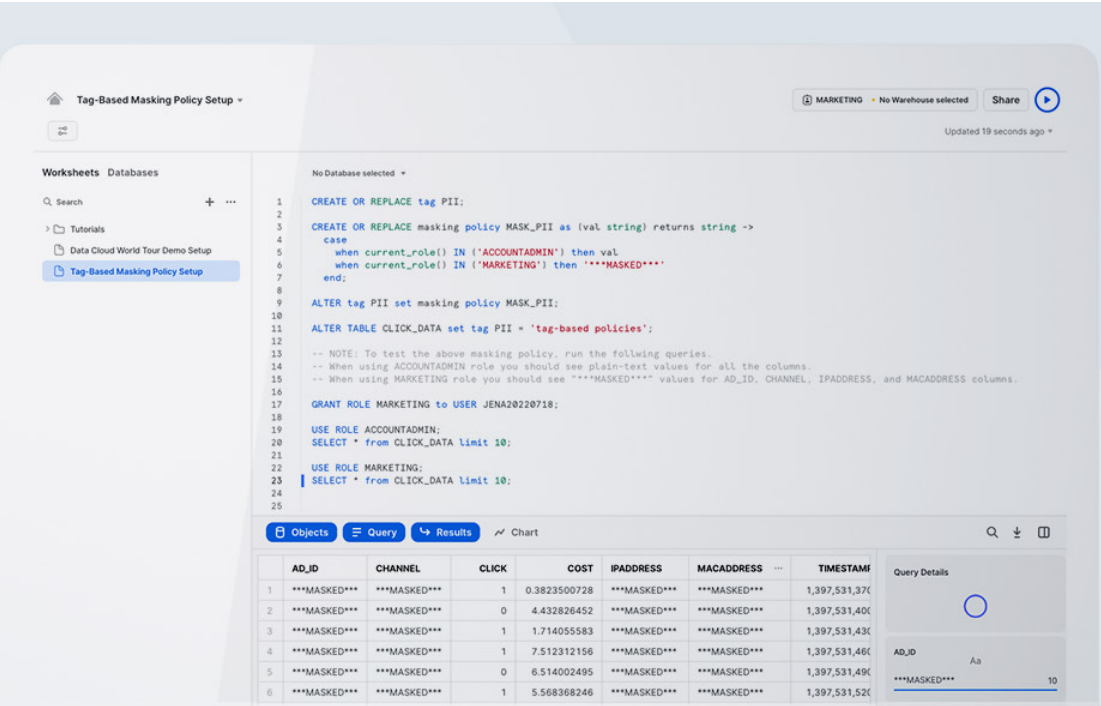
STEP 6: TEST DRIVE PROGRAMMATIC ACCESS AND SPARK/PYTHON

Once you have mastered executing queries using Snowsight, the web interface, and you have tested some standard BI tools, it is time to venture outside the database and test programmatic access. All requests to connect to Snowflake are made over a secure connection, and all executed queries are translated to REST calls. Using the Snowflake command-line interface (CLI) is a great place to start, but if you are more adventurous (or have code you would like to try with Snowflake), feel free to download Spark or Python connectors as well.

[Click here](#) for more details on programmatic access.

QUICK TIP

STEP 3
Be sure to test
business-critical
queries during
your trial.



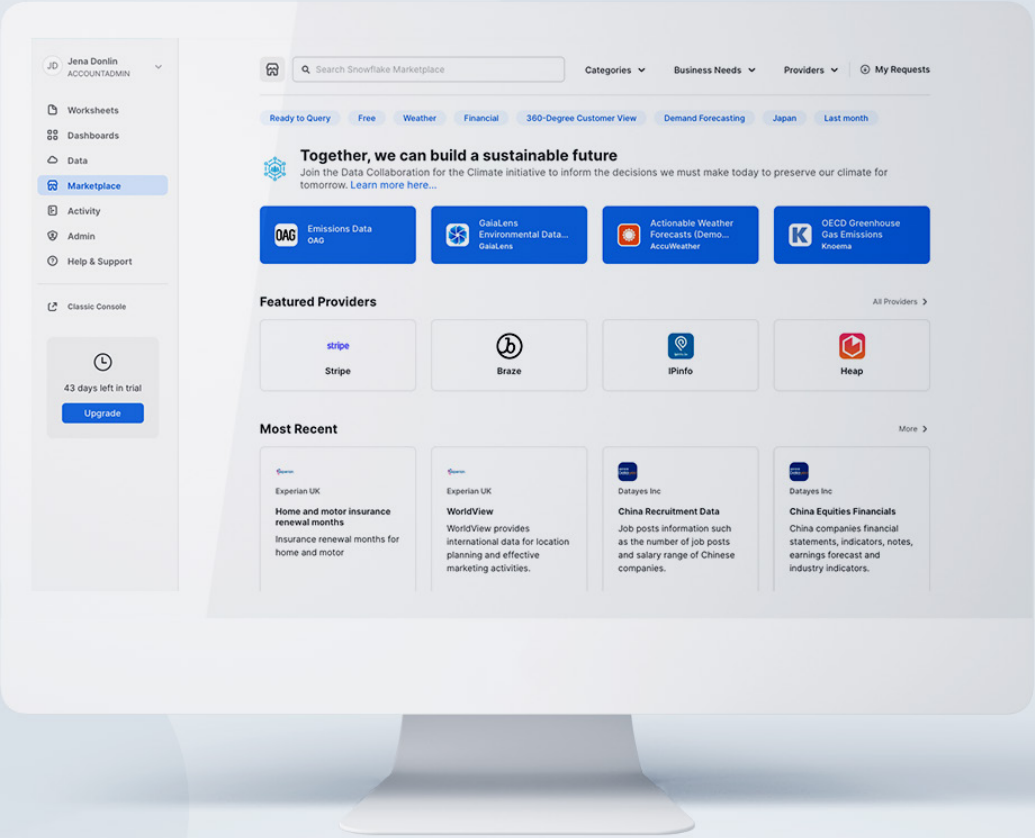
SHARING DATA

Before the Snowflake Data Cloud, if you wanted to share data, you had to use ETL processes to extract the data and send a copy of it to the other users. They in turn would have to create an ETL process to load the data into their system and keep it up to date. The process is expensive and time consuming to maintain, and results in stale data.

Snowflake Secure Data Sharing eliminates the hassle of copying and delivering data using traditional ETL processes or APIs, making data accessible with security and governance. No actual data is copied or transferred between accounts. Instead, all sharing is accomplished through Snowflake’s unique services layer and metadata store. It enables you to securely share selected objects with other Snowflake accounts. Tables, external tables, secure views, secure materialized views, and secure UDFs can all be shared with read-only access.

QUICK TIP: YOU’VE ALREADY USED DATA SHARING

You may not have realized it, but Snowflake shares its trial account database sample via Snowflake Secure Data Sharing.



QUICK TIP

Snowflake Marketplace contains third-party data sets to augment your own data.

EXPLORE DATA FROM SNOWFLAKE MARKETPLACE

Snowflake Marketplace gives you a one-stop shop for finding and accessing data sets that help you augment your own data. Browse Snowflake Marketplace listings and pick a free data set to explore. Add it to your trial account and try exploring and querying the data.

See the **Snowflake Marketplace Demo** for details about how to get started.

IF YOU HAVE MORE CREDITS TO BURN, TRY THESE IDEAS!

> UTILIZE SOME SEMI-STRUCTURED DATA

Snowflake is unique in that it can natively support **JSON** (and other semi-structured and unstructured data) alongside relational data. Most databases and data stores support only one format.

With Snowflake, you can choose to “flatten” nested objects into a relational table or store the objects and arrays in their native format within the VARIANT data type. JSON and other semi-structured data can be manipulated with ANSI-standard SQL, with the addition of dot notation.

> ACCESS HISTORICAL DATA

If you’ve been carefully managing your credits over time, you can use Snowflake **Time Travel** to access historical data. This includes data that has been changed or been deleted at any point within a defined period.

With **Time Travel**, you can:

- Restore data-related objects (tables, schemas, and databases) that may have been accidentally or intentionally deleted.
- Duplicate and back up data from key points in the past.
- Analyze data use and manipulation over specified periods of time.

[Click here](#) for more details on Time Travel.



NEXT STEPS

After you have used all of your credits and are satisfied with what Snowflake can provide, you can continue to use Snowflake with pay-as-you-go billing by simply adding a credit card to your account. You can also continue to test larger data sets and workloads, if needed, with pay-as-you-go billing.

If you still have questions or would like to inquire about a capacity contract, [contact Snowflake](#) to get answers to your questions.

Also, see the following resources:

- [Snowflake documentation](#)
- [Snowflake Inside the Data Cloud blog](#)
- [Snowflake on Medium](#)
- [Snowflake Community](#)
- [Snowflake Quickstart Guides](#)



ABOUT SNOWFLAKE

Snowflake enables every organization to mobilize their data with Snowflake's Data Cloud. Customers use the Data Cloud to unite siloed data, discover and securely share data, and execute diverse analytic workloads. Wherever data or users live, Snowflake delivers a single data experience that spans multiple clouds and geographies. Thousands of customers across many industries, including 506 of the 2021 Forbes Global 2000 (G2K) as of April 30, 2022, use Snowflake Data Cloud to power their businesses. Learn more at [snowflake.com](https://www.snowflake.com)



© 2022 Snowflake Inc. All rights reserved. Snowflake, the Snowflake logo, and all other Snowflake product, feature and service names mentioned herein are registered trademarks or trademarks of Snowflake Inc. in the United States and other countries. All other brand names or logos mentioned or used herein are for identification purposes only and may be the trademarks of their respective holder(s). Snowflake may not be associated with, or be sponsored or endorsed by, any such holder(s).