



# Enabling new products and services with Snowflake Data Sharing

## What can Snowflake Data Sharing do for data providers?

Snowflake Data Sharing is the most streamlined way to share data with your customers. Other methods for sharing data—such as, EDI, FTP and even e-mail—require data deconstruction, file transmission, copying, reconstruction, and duplicative storage. Even cloud file services require the copying and sending of files, an inherently manual, static and time consuming process.

With Snowflake's innovative cloud architecture, the world of sharing data is vastly simplified. Any Snowflake customer (data consumer) who wants to share data with someone else can give other organizations live access to their data over a secure connection. Rather than physically transferring the data to the consumer, the data remains in the provider's account and is visible and accessible to the consumer through SQL.

The costs for sharing data with Snowflake are minimal and straightforward. Data providers simply pay Snowflake for the data they store, and data consumers pay for the compute resources their queries consume. If the data provider chooses to charge the consumer for the data they use, Snowflake takes no fee or interest in the transaction.

With the freedom, simplicity and affordability of Snowflake Data Sharing, organizations everywhere have a host of new ways to optimize their most valuable resource: their data.

### POTENTIAL USE CASES FOR SNOWFLAKE DATA SHARING

#### Creating new data products

A data product is any dataset or stream that can be shared from one organization to another. By sharing data, the vast amounts of data resources that organizations are collecting can easily be turned into data products. Although there is potential to develop entirely new data products, in many cases companies are also monetizing their data as an add-on to products they've already developed.

Standalone data products are a popular and common way to monetize shared data. For instance, ratings companies sell their databases to advertisers eager to understand which TV shows are the most watched. Companies with existing data products like this can use Snowflake to reduce the manual effort and copying necessary with traditional data sharing tools. Companies who are new to data products can use data sharing to very quickly add entirely new business lines.

On the other hand, data products can also take the form of add-ons. For instance, the gaming company Playfab gets the majority of their revenue from their eponymous platform they sell to developers. However, they have also enabled Snowflake Data Sharing as an additional, optional add-on that their customers can purchase to see interaction and telemetry data that would ordinarily be difficult to collect and store. Playfab's customers gain the ability to get insight on their own games, and Playfab has a new (and easily enabled) stream of revenue.

### COMMERCIALIZATION CONSIDERATIONS

Snowflake does not participate in any monetization business model for data providers. Snowflake only makes money in the data sharing process from the data provider storing data, and the data consumer using compute resources to query it. As a data provider, you have complete ownership and freedom to choose how you commercialize your data.

If you do decide to commercialize your data and share it, there are a number of ways you can bill your consumers. Most simply, you can charge your consumers a flat one-time fee for a fixed dataset. For evolving datasets, or datasets that update over time, a subscription model would make more sense. For instance, you could charge customers a fixed fee per month of data. Expanding on that same concept, you could charge increasing fees for more rapidly updated data, enabling customers to buy up from monthly to weekly or daily subscriptions. Currently, data providers do not have access to consumer usage data, so that cannot be used to bill consumers.

### Adding value and enabling partners with data

Many organizations have no need to monetize their data—or they don't want to—but they still want to benefit from the opportunity to work more efficiently and collaboratively with other organizations. For instance, a trucking company may collect component failure data to share with the manufacturer to encourage the development of more reliable components. Or, a peer-to-peer lender may want to share their (anonymized) database with lenders to encourage analysis and further lending.

## WHAT CAN SNOWFLAKE DATA SHARING DO FOR DATA CONSUMERS?

Snowflake Data Sharing will also be an enormous benefit for your consumers, because the process of consuming shared data with existing methods is just as frustrating as the process for providing it. For example, with traditional data sharing, consumers must purchase their own data warehouse, which they will need to manage and maintain. Then, they need to manually reconstruct data as it is transmitted to them, a potentially agonizing task when data is frequently updated or changed.

Snowflake Data Sharing significantly improves the consumer experience. They can reduce the effort and expense of consuming shared data considerably because they don't need to pay only to store it: they only pay for the compute capacity they use to access the data in a provider's Snowflake account. Snowflake account. When data is added to a shared dataset, the changes are immediately visible to every consumer, making data updates a seamless and automated process.

## ENABLING SNOWFLAKE DATA SHARING FOR YOUR CUSTOMERS OR PARTNERS

Regardless of the data sharing business model you choose, Snowflake Data Sharing works the same. In order to enable and begin to share data, make the following key considerations:

1. What data are you sharing?
2. Who are your data consumers?
3. Are your data consumers already using Snowflake?
4. Is your account enabled for outbound shares?
5. Do your consumers need to access different subsets of data?

If your data consumers are not already using Snowflake, they will need [an account to get started](#). Once they create an account, you will need to put them in contact with your Snowflake Representative in order to enable data sharing on their account. If your data consumers are already active Snowflake customers, they will need to contact their Snowflake representative to enable data sharing on their account.

Technically, the process is simple and will be entirely handled by the Snowflake team. Once your data consumers have their accounts prepared and ready for data, be sure to obtain their account name from them; it's visible as the sub-domain in the URL when they login to Snowflake.



THE ACCOUNT IN THIS CASE IS "DEMO4"

By the same token, if your account has not been enabled for outbound shares, you will need to contact your Snowflake representative to enable the feature for you.

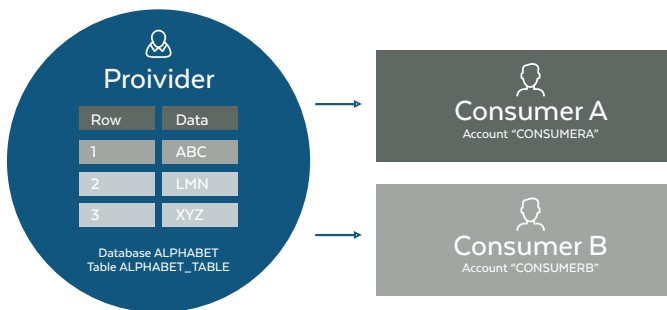
It's also very important to know whether you will be sharing your entire dataset with all of your consumers, or subsets of your data depending on the consumer. The process for sharing is slightly different depending on which of those scenarios you need to support.

## DATA SHARING SCENARIOS AND HOW TO ENABLE THEM

Snowflake Data Sharing is a powerful tool that can enable a host of different scenarios of varying complexity. It's important to identify the scenario that most closely matches your use case so you can enable sharing in the way that works best for you.

### Sharing the same data with one or multiple consumers

In this situation, the need is simple. You have a table or view (or many of them) that you would like to share with one or many consumers. For instance, if a federal agency wanted to share its entire database with all 50 state agencies, it would use this method.



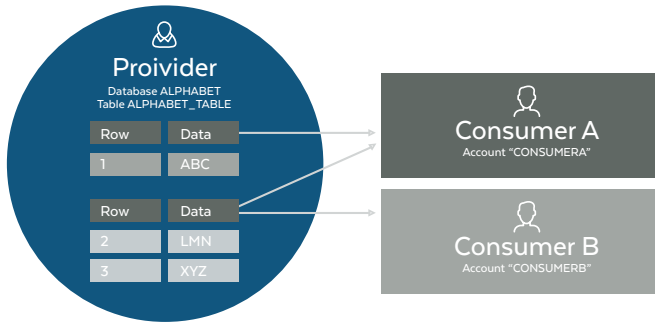
In the previous example, the provider is sharing the entire table ALPHABET\_TABLE within the database ALPHABET with consumer A and consumer B in similarly named accounts. The process is relatively straightforward. First, you will need to create a share, which is a proprietary database object in Snowflake that serves as the mechanism for conferring permissions and data access to your consumers. To complete the process, you will need to add permissions for all related database objects to the share, and then add the consumer accounts to the share.

```
...  
  
//1 Create the share  
  
Create share MY_ORGS_DATA;  
  
//2 Grant usage on the share to all  
pertinent database objects  
  
Grant usage on database ALPHABET to share  
MY_ORGS_DATA;  
Grant usage on schema ALPHABET_SCHEMA to  
share MY_ORGS_DATA;  
Grant select on table ALPHABET_TABLE to  
share MY_ORGS_DATA;  
  
//3 Add the accounts to the share  
  
alter share MY_ORGS_DATA add accounts  
= CONSUMERA, CONSUMERB;
```

The process would be similar for a secure view and could be extended to multiple views or tables simply by granting access on those specific objects to the share.

### Sharing different subsets of data with multiple consumers

The second scenario is more complicated and could have multiple iterations, but the underlying need is to share different subsets of data with multiple different consumers. For instance, a car manufacturer may wish to share future production by state with its dealers. However, some dealers may be confined to a single state, while others may span across many states. Using this method, the data provider could share the pertinent data with each dealer, regardless of their location(s).



In this scenario the same data as the previous example is used, but it is broken up into two “groups”. Group 2 contains rows 2 and 3 and is visible to all consumers. Group 1 contains row 1 and is visible to only the first group.

Facilitating this kind of sharing is not difficult, but it does require a couple of extra steps over the simplified data sharing example above. First, the data provider will need to create an `account_mapping` table which matches the consumer account names to the “group” that each consumer has access to. Make sure that the account names are uppercase, as Snowflake will not recognize account names in lowercase. The account mapping table will look similar to this:

Table: ALPHABET\_TABLE

Row	Data	Group
1	ABC	1
2	LMN	2
3	XYZ	2

Table: ACCOUNT\_MAPPING

Group	Account_name
1	CONSUMERA
2	CONSUMERA
2	CONSUMERB

Second, the data provider will need to create a secure view. A secure view functions in a similar fashion as normal views, in that they allow the result of a query to be accessed as if it were a table. However, some of the internal optimizations for normal views require access to the underlying data in the base tables for the view. This access may allow data that is hidden from users of the view to be exposed through user code, such as user-defined functions, or other programmatic methods. Secure views do not utilize these optimizations, ensuring that users have no access to the underlying data.

After creating a secure view, the provider will need to reference the `CURRENT_ACCOUNT()` function in Snowflake to align the group permissions in the account mapping table to the data that each account can view in the underlying table. In other words, the secure view

looks at the account that is accessing data, matches that account name to the groups in the `account_mapping` table, and then returns the rows that the account has permissions to view.

The rest of the steps are the same as in the previous example: create the share, grant privileges on the objects that are being shared and add each consumer account to the share.

```

...

//1 Create a secure view

Create secure view ALPHABET_SECURE as
select row, data, group from ALPHABET_TABLE
a join ACCOUNT_MAPPING am on a.group =
am.group

where am.account = CURRENT_ACCOUNT();

//2 Create the share

Create share MY_ORGS_DATA;

//3 Grant usage on the share to all
pertinent database objects

Grant usage on database ALPHABET to share
MY_ORGS_DATA;
Grant usage on schema ALPHABET_SCHEMA to
share MY_ORGS_DATA;
Grant select on table ALPHABET_TABLE to
share MY_ORGS_DATA;

//4 Add the accounts to the share

alter share MY_ORGS_DATA add accounts =
CONSUMERA, CONSUMERB;

```

## ENSURING CONSUMERS ARE SUCCESSFUL WITH INBOUND DATA SHARES

Although it's incredibly easy to enable inbound data shares for consumers, it's also vital to support them through the beginning of the process to make sure they have a good experience. The following questions provide a framework for that discussion:

### Are your consumers using Snowflake?

First of all, it's important to make sure that your consumers are comfortable using Snowflake in the first place. It doesn't require much at all: they can use

standard SQL, and there's no tuning, optimization or manual management to worry about. There are a number of resources to help if your consumers are new to Snowflake (or databases in general), including a detailed documentation site with [Getting Started tutorials](#) and guidelines.

### Are they comfortable enabling data sharing on their own?

Although SQL is a common skill, it is possible that some customers or partners may want an even more straightforward method for accessing the data that is shared with them. For instance, a common request would be to access data through Looker or Tableau (both Snowflake partners). There are many ways to enable this, including a templated method from Looker called Looker Blocks. However, you will still need to help your consumer get started with their Snowflake account so they can create a database from the share in order to facilitate the connection with their BI tool.

The following process can be either shared with the customer to perform on their own or used to assist them manually.

### HOW TO HELP YOUR CONSUMERS USE THE DATA YOU HAVE SHARED WITH THEM

The process of consuming data through a data share is exceedingly simple and functions exactly the same regardless of the method the provider used to share. However, it is important to follow the data consumption process precisely since each step is dependent on the previous one.

1. Enable Snowflake inbound data sharing (if not already enabled).
2. View shares.
3. Create a database with the inbound share.
4. Grant imported privileges to appropriate roles.

If the account isn't already enabled for inbound data sharing, the consumer will need to reach out to their Snowflake Representative to have it modified. Once the account is enabled, the consumer can view any inbound data share by using the following SQL:

```
...  
show shares;
```

The result in the Snowflake worksheet will show both inbound and outbound shares, along with additional information on the database that the share is based upon.

The shares themselves are referred to with the following syntax:

```
...  
[provider account].[share name]
```

In order to put the share to use, the consumer needs only to create a database with the inbound share and apply the imported privileges to pertinent roles within Snowflake. The following would enable anyone in the consumer's Snowflake account to view the data, but could be modified [to enable only specific roles](#).

```
...  
  
//1 Create a database from the share  
  
create or replace database ALPHABET  
from share [PROVIDER_ACCOUNT].MY_ORGS_DATA;  
  
//2 Query to view the data in the shared  
table  
  
select * from ALPHABET;  
  
//3 Grant permissions to others  
  
grant imported privileges on database  
ALPHABET to role public;
```

## WHAT CAN CONSUMERS DO WITH THEIR INBOUND DATA SHARE?

At this point, the data will be visible to the consumer account and can be queried with standard SQL. However, there are a couple of important differences between shared data and data that has been natively loaded into Snowflake.

1. The consumer cannot alter the underlying data in a share. On the same token, when the data is updated in the provider account, the changes will immediately be visible in the consumer account. In other words, the data is live and requires no transmission or copying for the consumer to access as it updates.
2. If a consumer wishes to alter a database, view or table from an inbound data share, they may select and copy it into a new database. However, this new database will not be updated when data changes in the provider account.

3. Consumers cannot share an inbound share with any other Snowflake account. In other words, you cannot share a share.

Otherwise, consumers are free to query, permission, view and use their inbound data exactly as they would any other database in Snowflake. What's more, they can combine and utilize their inbound data shares with their existing corporate data without painful extract, load, transform (ELT) processes.

## SUMMARY

Effectively sharing data is easier than ever with Snowflake Data Sharing. Any data provider can create new data products for their customers or partners with easy-to-understand SQL and a Snowflake account. The opportunities for expanded revenue and increased value are nearly endless, and they only become more numerous as the volume of data increases.

If you are interested in getting started with data sharing today, contact your Snowflake Representative. If you are new to Snowflake, start a new [On-Demand account](#) and claim your \$400 in free usage on Snowflake today.



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. Snowflake is headquartered in Silicon Valley and can be found online at [snowflake.net](https://www.snowflake.net).