



FOCUS ON YOUR DATA, NOT YOUR DATA WAREHOUSE

To participate in the virtual hands-on lab please login to your account.
Sign up for a Free Snowflake Trial: <https://signup.snowflake.com>

HANDS-ON LAB SET UP

- Sign-in to your free trial account. URL looks something like:
<https://XY12345.snowflakecomputing.com>
- If you have not already signed up for a free trial account, sign up here:
 - <https://signup.snowflake.com>
 - Any provider/region/edition will work. But we suggest region closest to you and Enterprise Edition
- Review the guide [Getting Started with Snowflake - Zero to Snowflake](#) at guides.snowflake.com
 - Load the SQL code from Step 2 of the Guide into a Snowflake worksheet



MEET TODAY'S HANDS-ON LAB LEADERS



Brock Cooper
Sales Engineer



Blake Snyder
Account Executive



WHAT WE'LL COVER TODAY

Snowflake Introduction

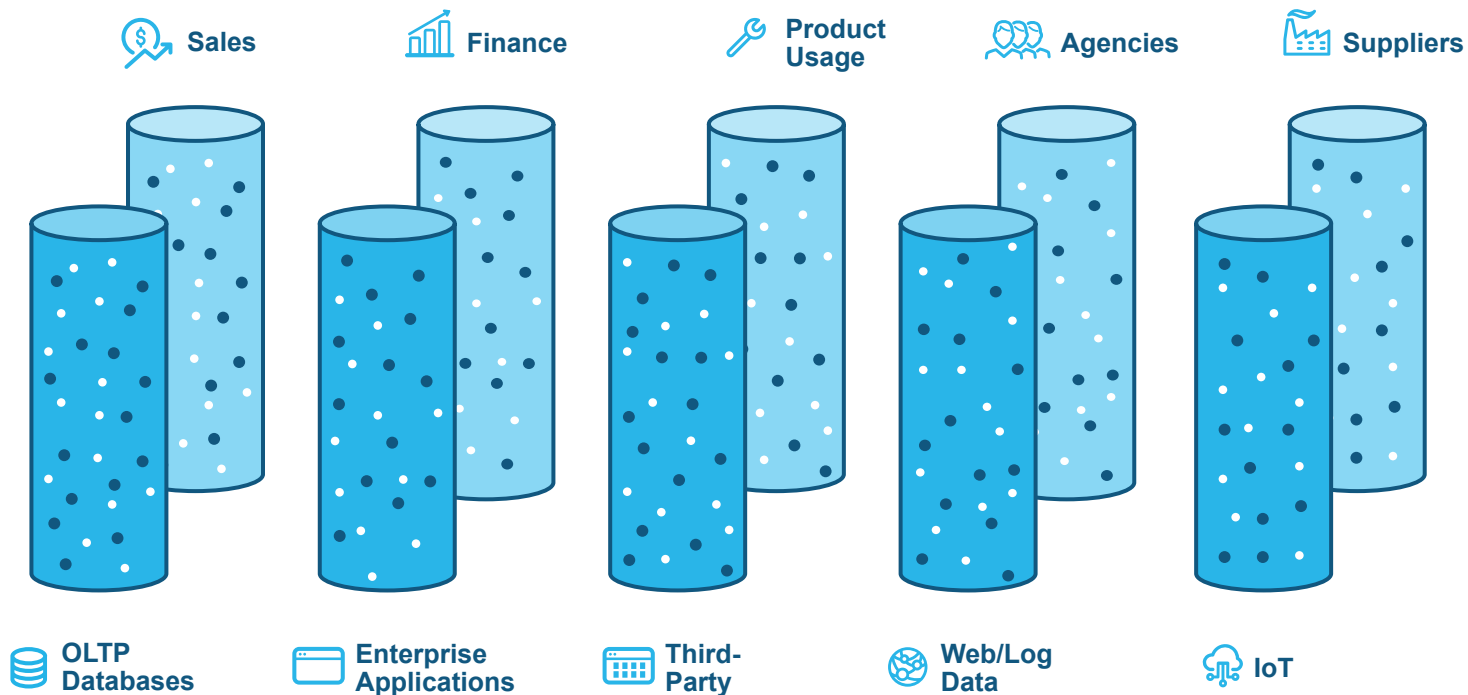
Hands-on Lab

- > Create stages, databases, tables, views, and warehouses
- > Load structured and semi-structured data
- > Query data including joins between tables
- > Clone objects
- > Undo user errors
- > Create roles and users, and grant them privileges


Q&A




DATA SILOS PREVENT VALUE REALIZATION



THE DATA CLOUD IS A GLOBAL NETWORK



THE DATA CLOUD IS A GLOBAL NETWORK



1

ACCESS

2

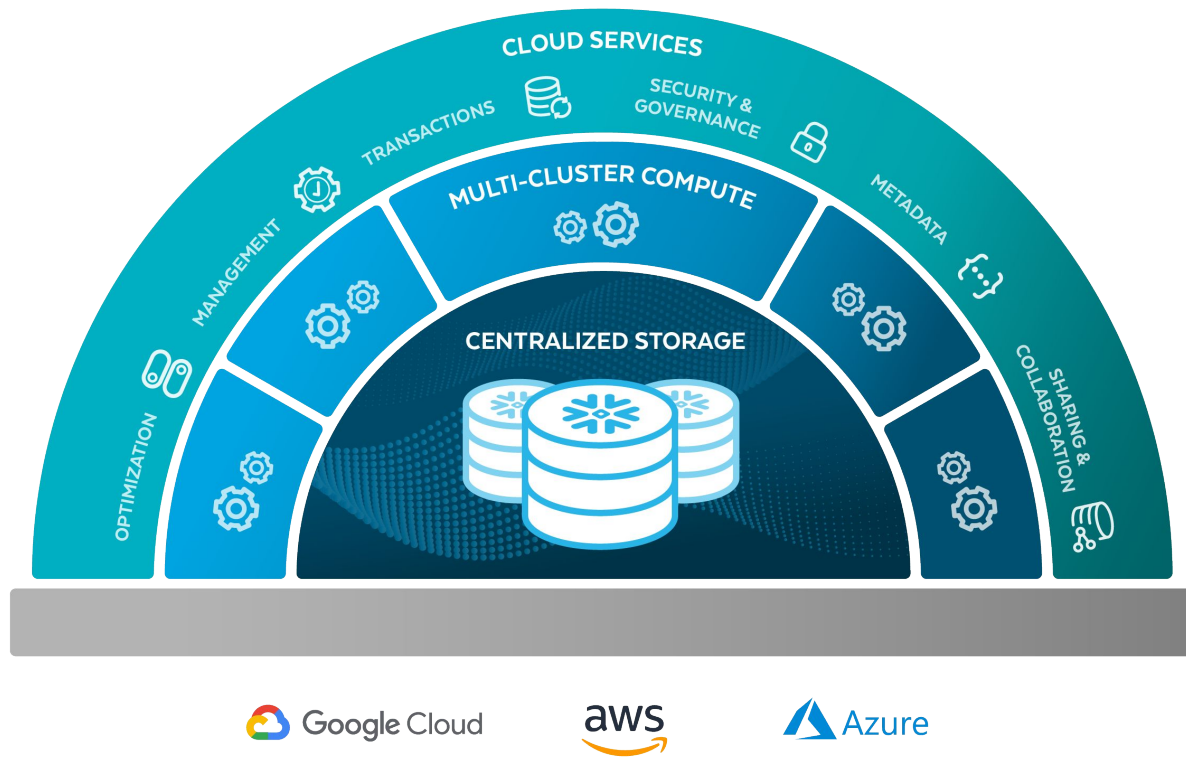
GOVERNANCE

3

ACTION



SNOWFLAKE ARCHITECTURE



**LET'S DIVE INTO
SNOWFLAKE!**



Citibike Data



Station A

Name
Location (Lat/Long)
Type
Capacity
Region



Trip

Start time
End time
Start Station
End Station
Bike Id



Station B

Name
Location (Lat/Long)
Type
Capacity
Region

Data Profile

Date Range: 2013 - 2020

Number of trips: 100M+

Number of stations: 1200

Questions from the data

Where do they ride? When do they ride?

How many times do they ride? From where and to where do they ride?

What was the temperature when they rode?



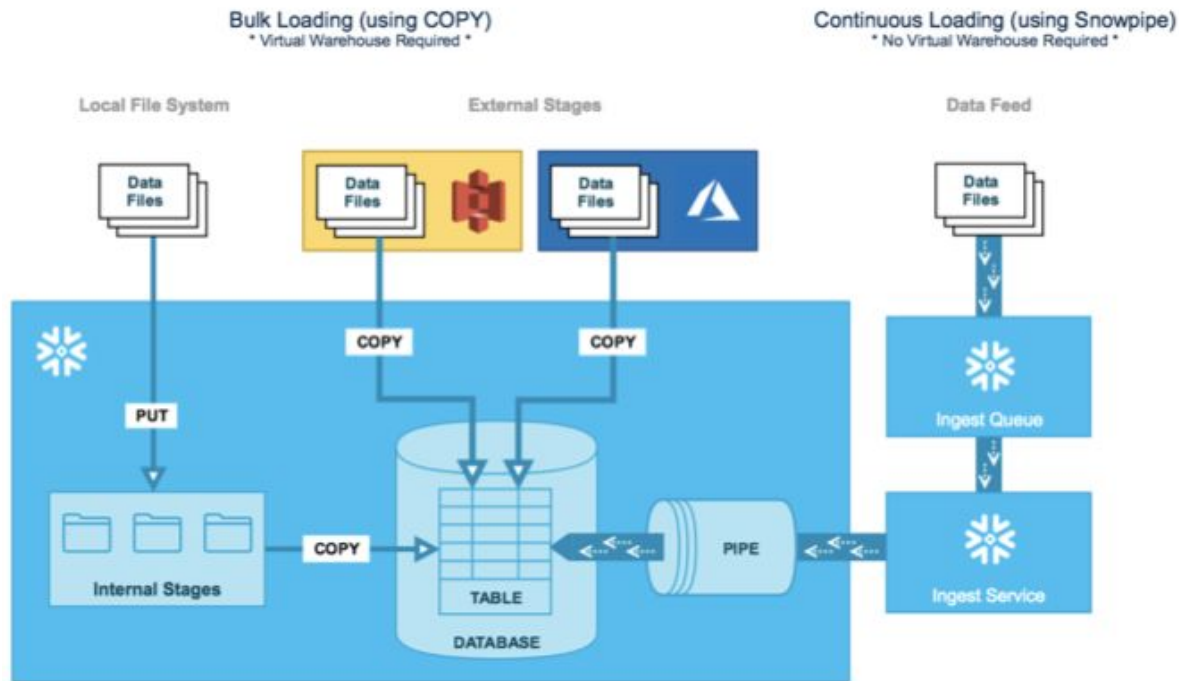
HOW TO LOAD DATA INTO SNOWFLAKE

Bulk Loading

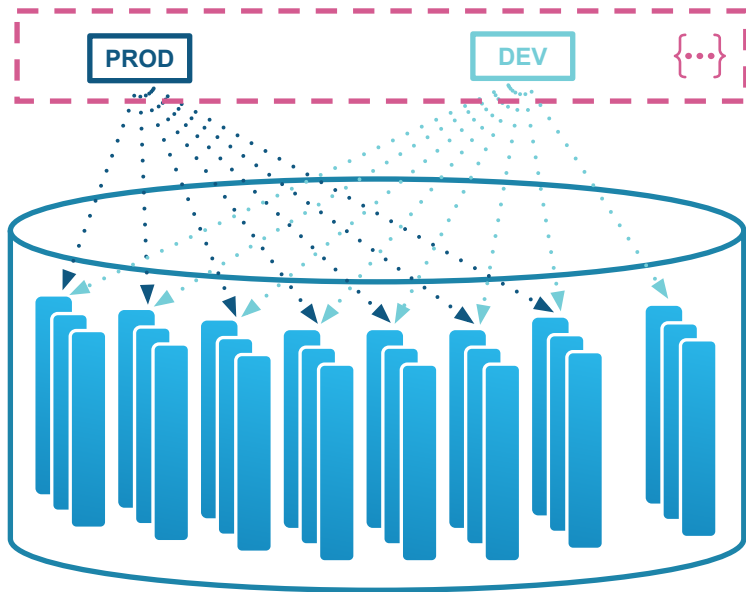
- COPY Command
- User-managed compute resource

Continuous Loading

- Snowpipe
- Snowflake-managed compute resource



ZERO-COPY CLONING



The Metadata layer keeps track of every micro-partition file in every customer database.

Creating a DEV environment usually means copying the PROD database

Limited to subset of full Prod

Up to 2x storage requirement

Periodic refreshes

Snowflake Zero-Copy Clones

Simply “point” to the same files

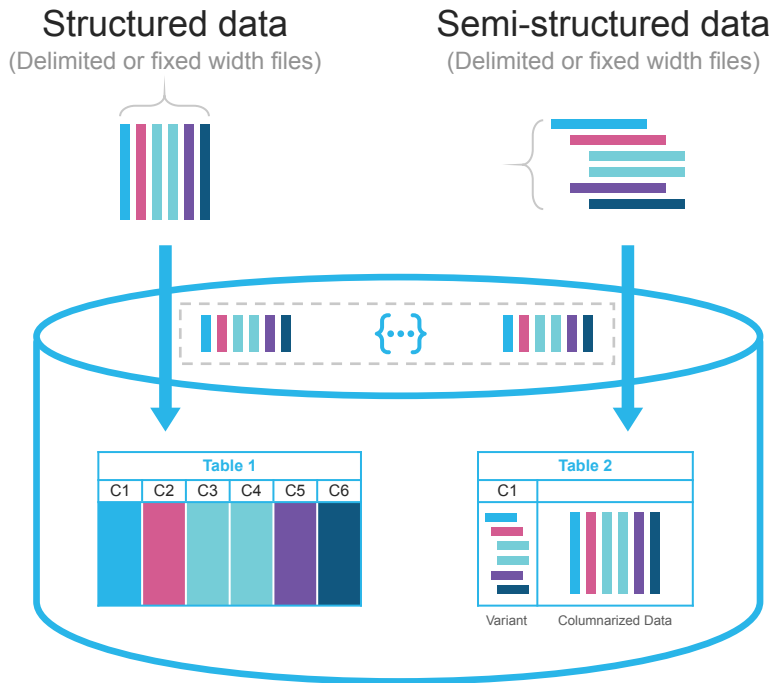
Consumes zero additional storage

Changes to either DB are isolated



SEMI-STRUCTURED DATA

Native support for JSON, XML, Avro, Parquet, OR



Structured formats (CSV, TSV, ...)

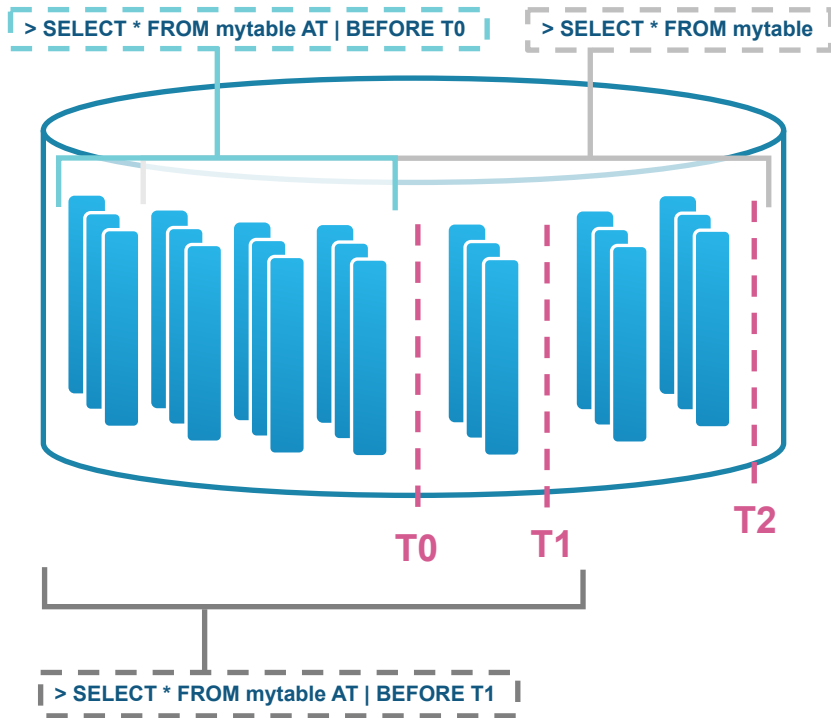
- Strongly-typed “columns”
- Typically map 1:1 to table columns
- Ingestion process generates important metadata

Semi-structured formats (JSON, XML, ...)

- Traditional DBs require complex transforms to “flatten” data into structures.
- Snowflake has a VARIANT datatype
- Stores original document as-is
- During ingestion, data is columnized and metadata collected
- SQL syntax is a simple dot notation



TIME TRAVEL



T0 – Initial state of database

T1 – `update myTable set
colX = Y where...`

T2 – ELT job loads new data

**Previous versions of data
automatically retained**

`AT | BEFORE [timestamp | statement | offset]`

`CLONE AT | BEFORE` to recreate a prior version

`UNDROP` recovers from accidental deletion

Accessed via SQL extensions

`AT | BEFORE [timestamp | statement | offset]`

`CLONE AT | BEFORE` to recreate a prior version

`UNDROP` recovers from accidental deletion



WHAT WE ACCOMPLISHED TODAY

1. **Navigating the UI**
2. **Creating databases and data warehouses**
3. **Loading & querying structured and semi-structured data**
4. **Zero copy cloning**
5. **Undoing user errors**
6. **Role based access control**
7. **Data Marketplace**





THANK YOU

snowflake.com/contact

snowflake.com/virtual-hands-on-lab/