SNOWFLAKE DATA CLOUD DEMO

Unite Your Siloed Data. Access a World of Data.

Meet today's presenters

David Kincaid
Senior Sales Engineer
Snowflake

Joseph Wu
Senior Sales Engineer
Snowflake

Details

- The webinar will be shared with the participants within the next couple of days.
- If you have any questions, please use the Q&A functionality and provide adequate context.
- If you have any audio issue, refresh your browser.
- Let's get started ...

Innovation Journey to the Data Cloud

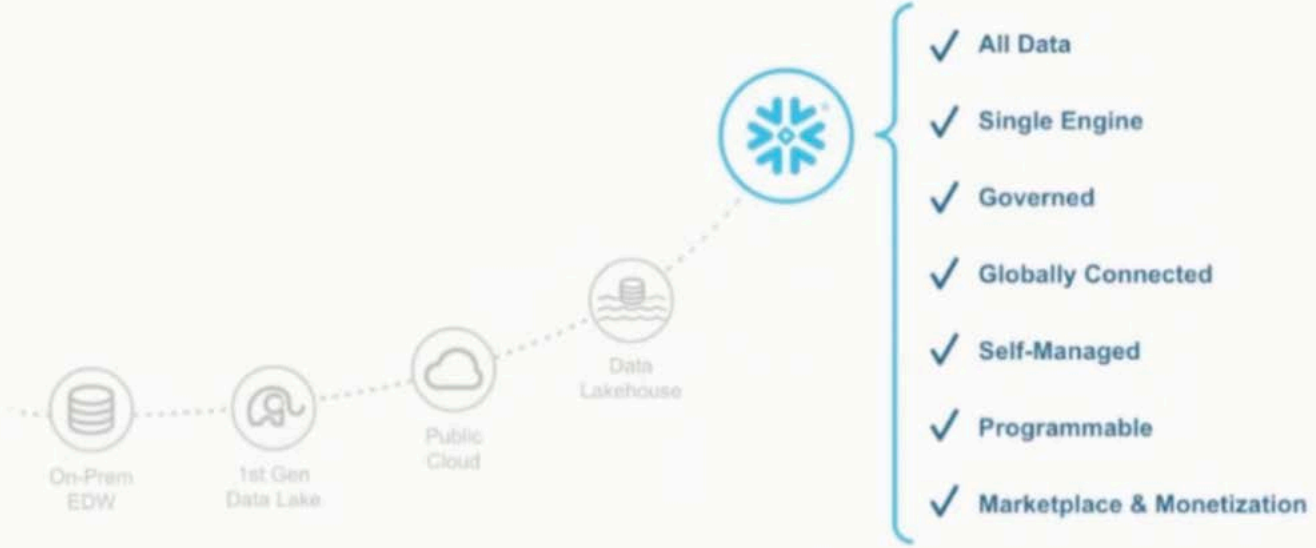

Modern Platform Requirements

Fast For Any Workload

Run virtually any number or type of job across users and data volumes quickly and reliably.

It Just Works

Replace manual with automated to operate at scale, optimize costs, and minimize downtime.

Connected to What Matters

Extend access and collaboration across teams, workloads, clouds, and data, seamlessly and securely.

SNOWFLAKE PLATFORM ARCHITECTURE

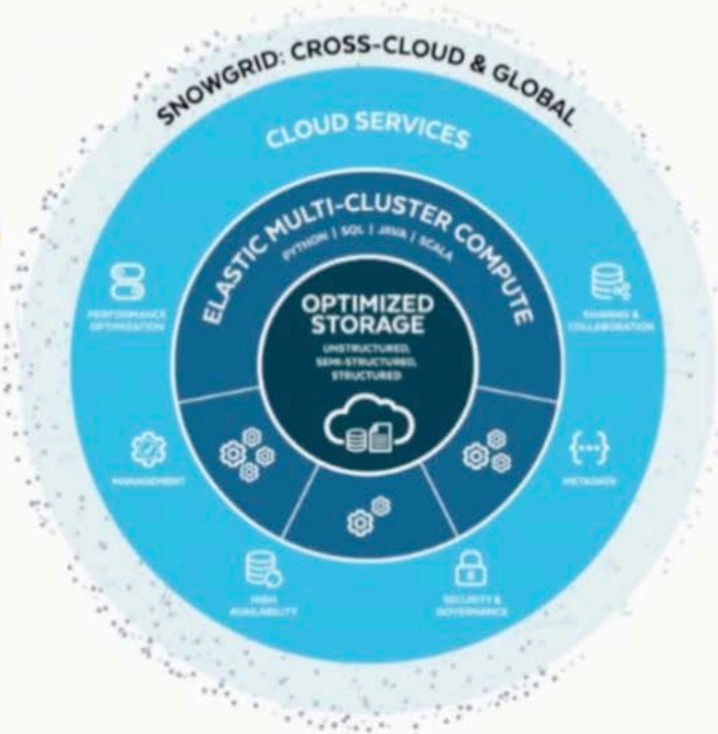

aws

Google Cloud

Optimized Storage

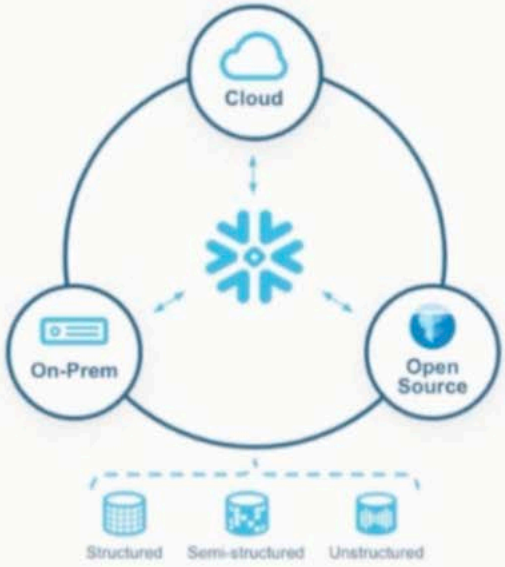

Unsiloed access to your data

Unstructured, semi-structured, and structured data together with near-infinite scale.

Easily manage data at scale

Fast and efficient access, optimized compression, and secure data - all automated.

Flexibility & interoperability

Work with data on-premises* or in open table formats* to remove lock-in and adapt to new data patterns.

* in private preview

Elastic Multi-Cluster Compute

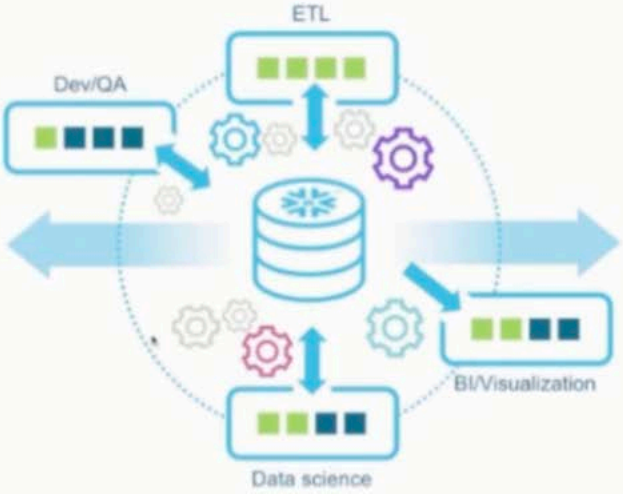

One engine for every workload

Simplify your architecture. Power complex pipelines, analytics, data science, interactive applications, and more.

Leading performance and concurrency

Fast, reliable performance for virtually all users and jobs with no tuning or contention.

Accessible & programmable

Work in SQL, Python, or Java, and run your preferred tools and libraries directly with Snowpark - without moving data.

Code the Same Way, Execute Faster With Snowpark

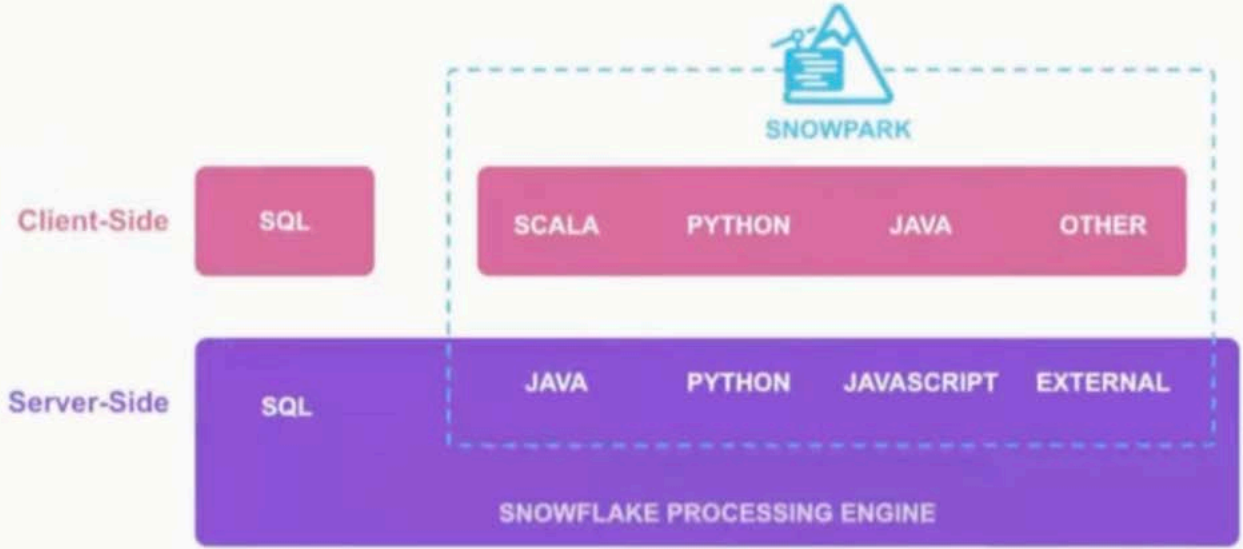

Cloud Services

Self-managed

Automate encryption, access controls, availability, tuning, maintenance, and more to keep operations simple and smooth.

Transparent improvements

Continually benefit from the latest performance enhancements and economics - no action required.

Optimized resources & costs

Only pay for what you used and get full visibility and cost governance controls to right-size costs.

Unified Governance

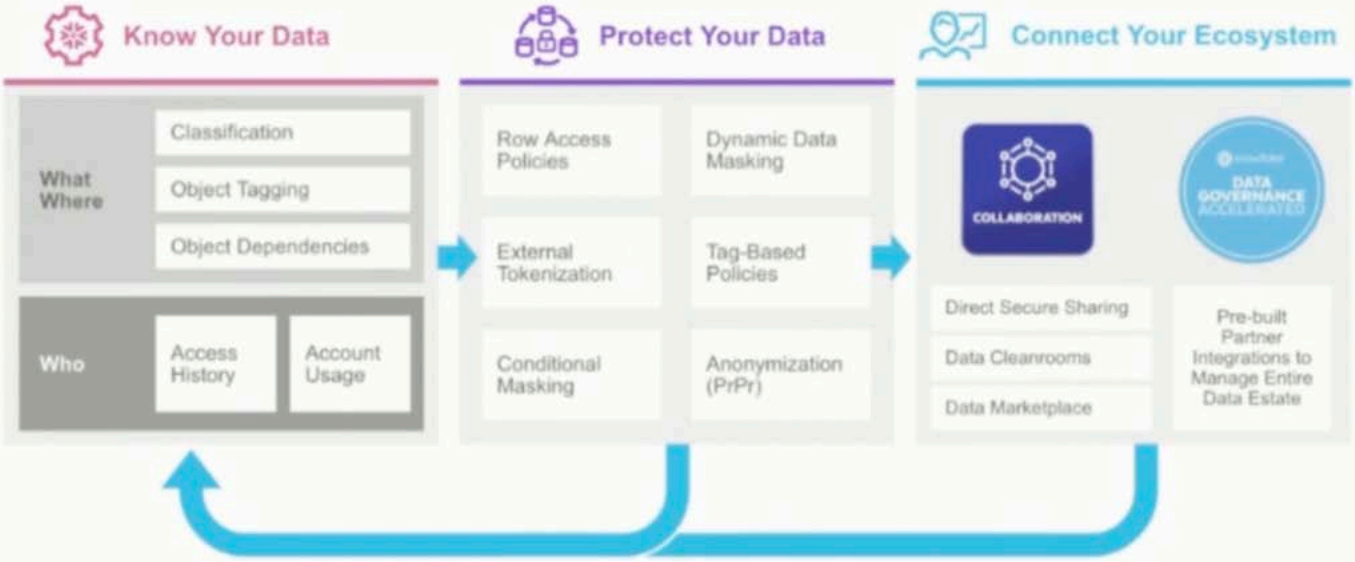

Data Cloud Growth

April 2023

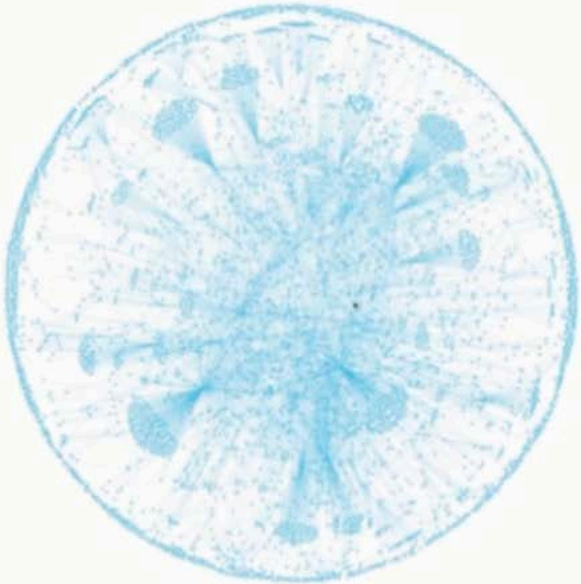

April 2023

* Visualization based on actual Data Cloud sharing activity as of April 30, 2023 and April 30, 2023 respectively

One Platform

Powering Many Workloads

Discover, access and monetize live data, services and apps in the Data Cloud

Build simple, reliable data pipelines at scale in the language of your choice

Protect your enterprise with near-unlimited visibility, unified data, and powerful analytics

Accelerate your ML workflow with fast access and elastically scalable processing

Build data-intensive applications without operational burden

Accelerate analytics for users and queries with leading price / performance and no complexity

Deploy flexible architectural patterns with governance and optimized storage at scale

Delivers a modern approach to working with transactional and analytical data together

Live Demo

WHO IS TASTY BYTES?

ABOUT US: Global food truck network, localized menu options, 15 countries, 30 major cities, and 15 core brands.

OUR MISSION

We serve to give people unique food options with high quality items in a safe, convenient and cost effective way. We ensure that the ingredients used are of the highest quality from mostly local food vendors to make sure our success has a positive impact on community partners.

OUR VISION

To become the largest food truck network in the world by 2027 that has sustainable profitability with a zero carbon footprint future that our team, customers, and communities are proud of supporting.

LOCATIONS SERVED

- USA: San Mateo, Denver, Seattle, Boston, New York City
- Canada: Toronto, Vancouver, Montreal
- United Kingdom: London, Manchester
- France: Paris, Nice
- Poland: Warsaw, Krakow
- India: Mumbai, Delhi
- Japan: Tokyo
- South Korea: Seoul
- Australia: Sydney, Melbourne

CURRENT STATE & FUTURE GOALS

	current		future
Trucks	450	*20% YoY growth over 5 years*	1120
Sales	$105M/year	*25% YoY growth over 5 years*	$320M/year
NPS	3	*Improvement by YE 2023*	40

frostbyte.

WAREHOUSE MANAGEMENT

Warehouse != Data Storage
Database == Data Storage
Warehouse == Compute

Dimensions of Scaling

ACROSS

- Many competing workloads
- Resource contention
- Isolate on separate warehouses

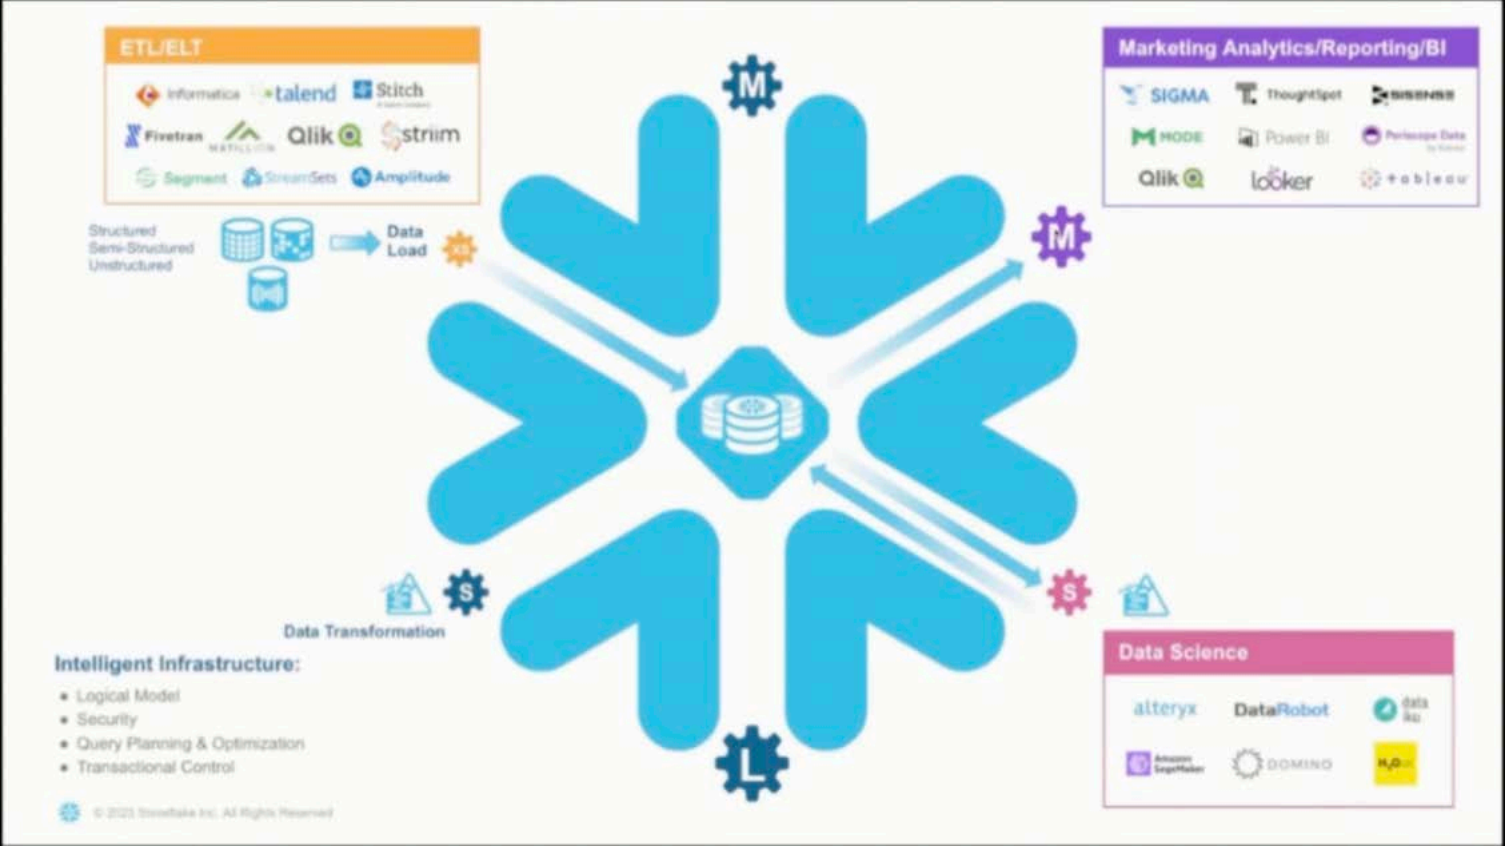

Dimensions of Scaling

UP

- Single query performance
- More data, more complex queries
- Add more resources to the cluster

ACROSS

- Many competing workloads
- Resource contention
- Isolate on separate warehouses

ETL/ELT

Informatica · talend · Stitch · Fivetran · MATILLION · Qlik · striim · Segment · StreamSets · Amplitude

Structured
Semi-Structured
Unstructured

Data Load

Marketing Analytics/Reporting/BI

SIGMA · ThoughtSpot · SISENSE · MODE · Power BI · Periscope Data · Qlik · looker · +ableau

Data Transformation

Data Science

alteryx · DataRobot · data iku · Amazon SageMaker · DOMINO · H2O.ai

Intelligent Infrastructure:

- Logical Model
- Security
- Query Planning & Optimization
- Transactional Control

Scale Up – Loading 1BN Records

- Doubling the number of servers halves the run time
- But you pay per-server, per second of compute
- So you get your answer

8X FASTER FOR THE SAME COST

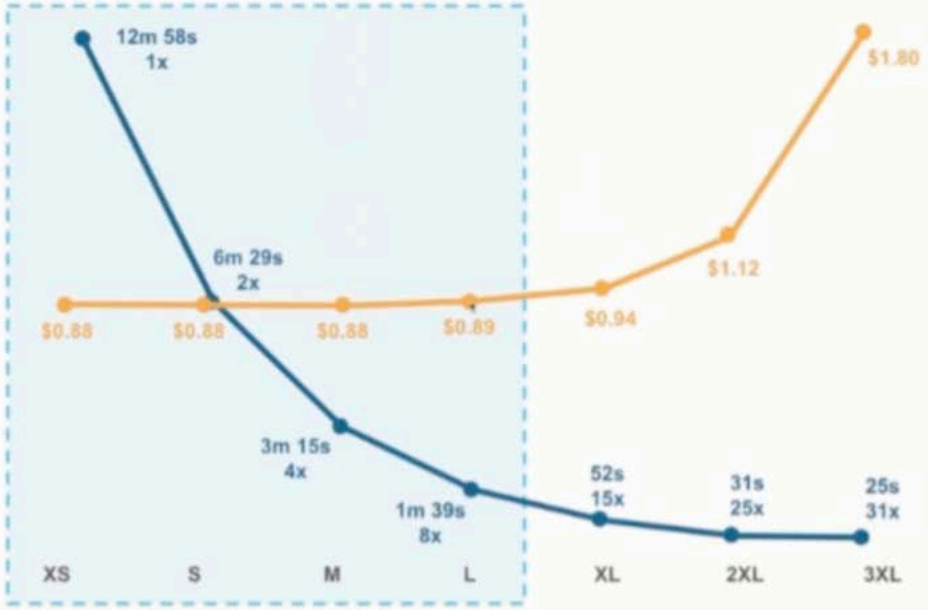

3 Dimensions of Scaling

UP

- Single query performance
- More data, more complex queries
- Add more servers to the cluster

ACROSS

- Many competing workloads
- Resource contention
- Isolate on separate warehouses

Out

- More users
- More queries simultaneously
- Spin up more clusters

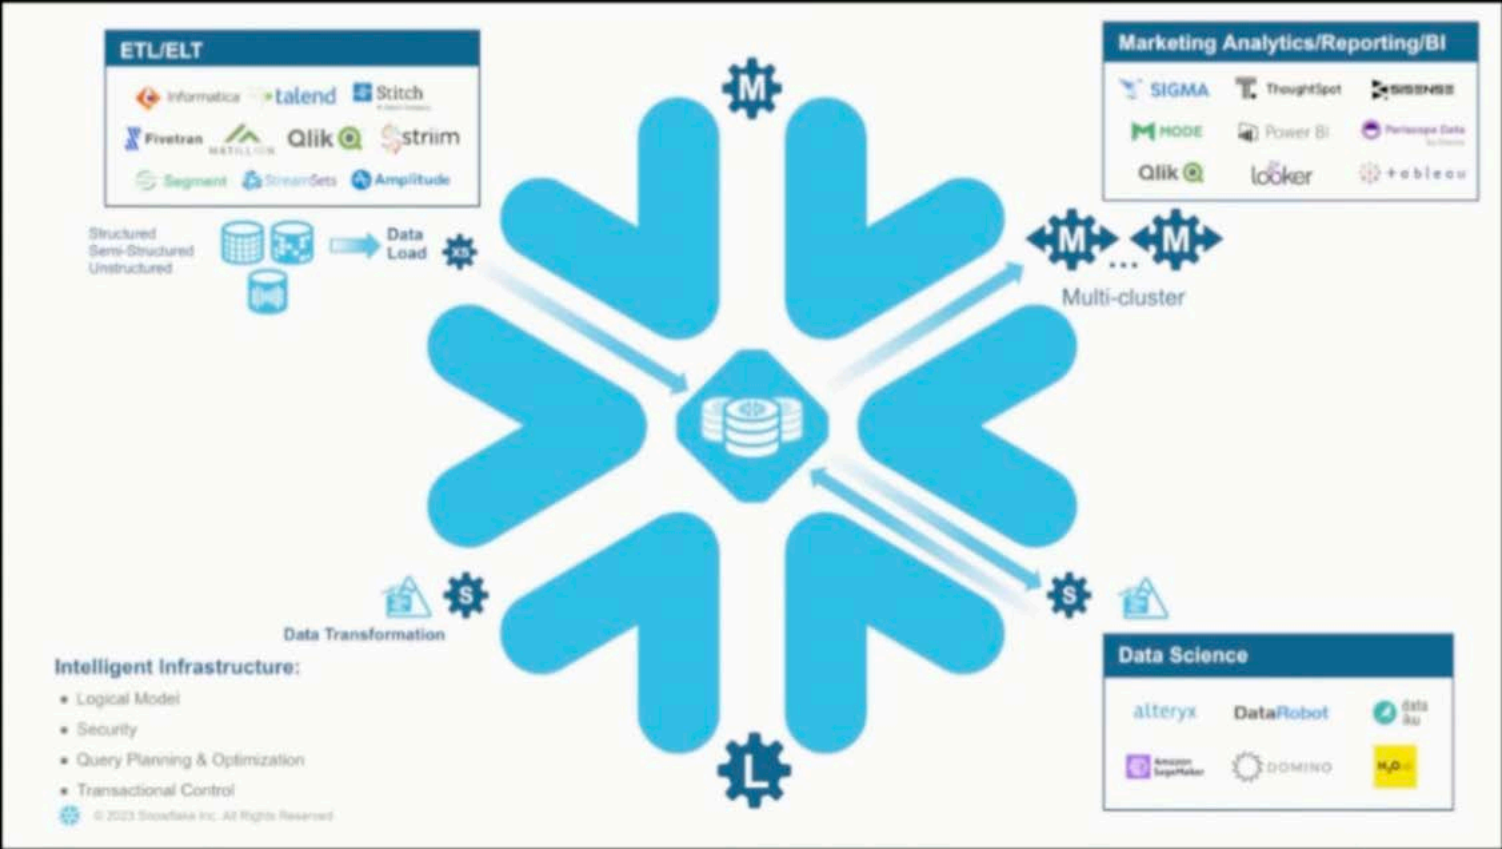

Warehouses/Compute Demo

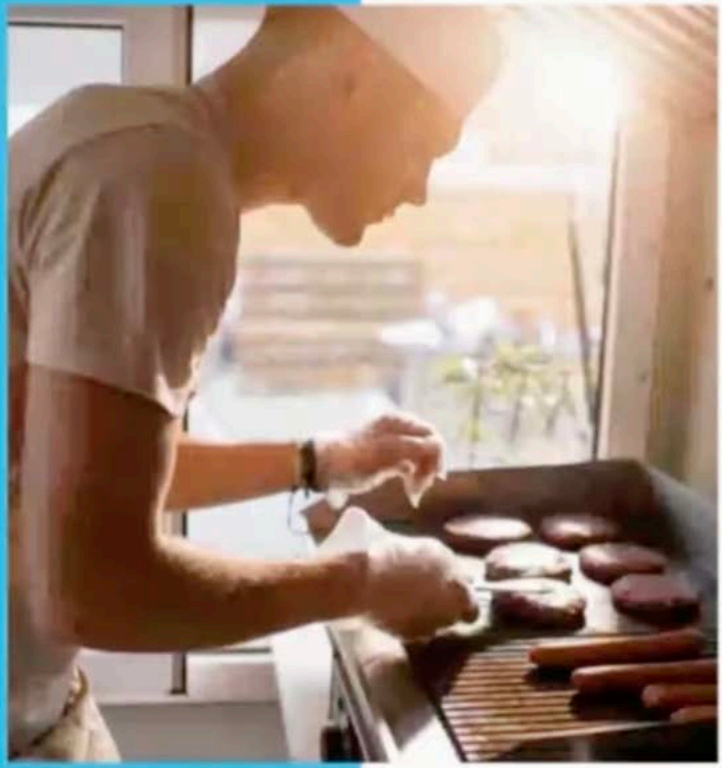

TRANSFORMATION

Zero-Copy Cloning

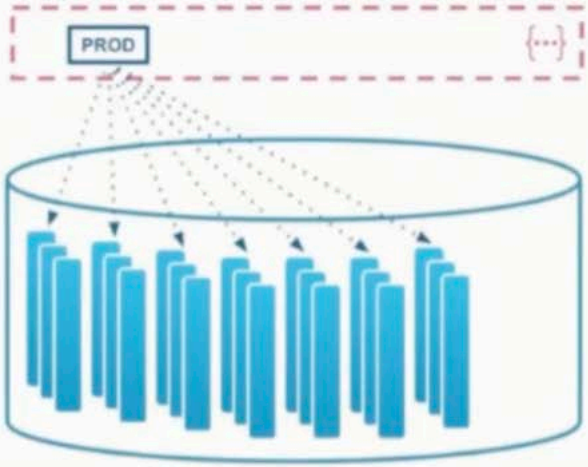

The Metadata layer keeps track of every micro-partition file in every customer database.

Creating a DEV environment usually means copying the PROD database

Limited to subset of full Prod

Up to 2x storage requirement

Periodic refreshes

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

Time Travel

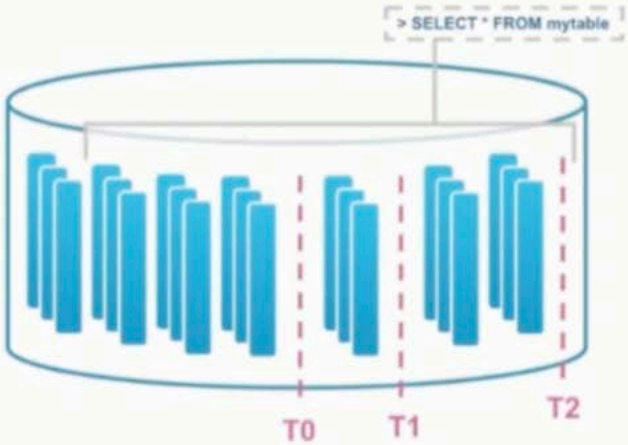

T0 – Initial state of database

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

Transformation Demo

frostbyte.

COLLABORATION

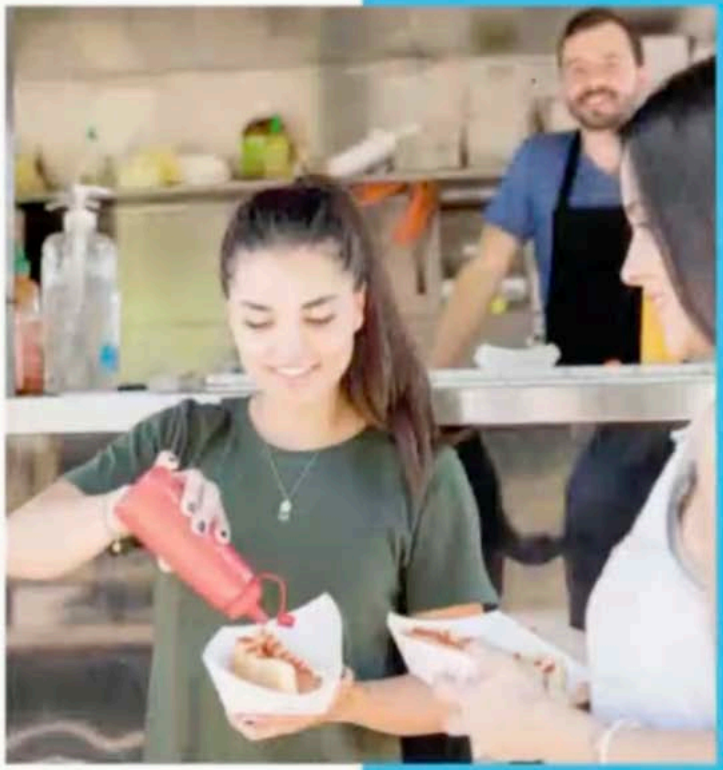

Collaboration

- Any Format
- Near-Unlimited Scale
- Sharing Without Copying or Moving

All of Your Organization's Data, on One Platform

Your Ecosystem - Partners, Suppliers, Customers

Third-Parties - Industry Datasets, Data Services, Applications

Code the Same Way, Execute Faster With Snowpark

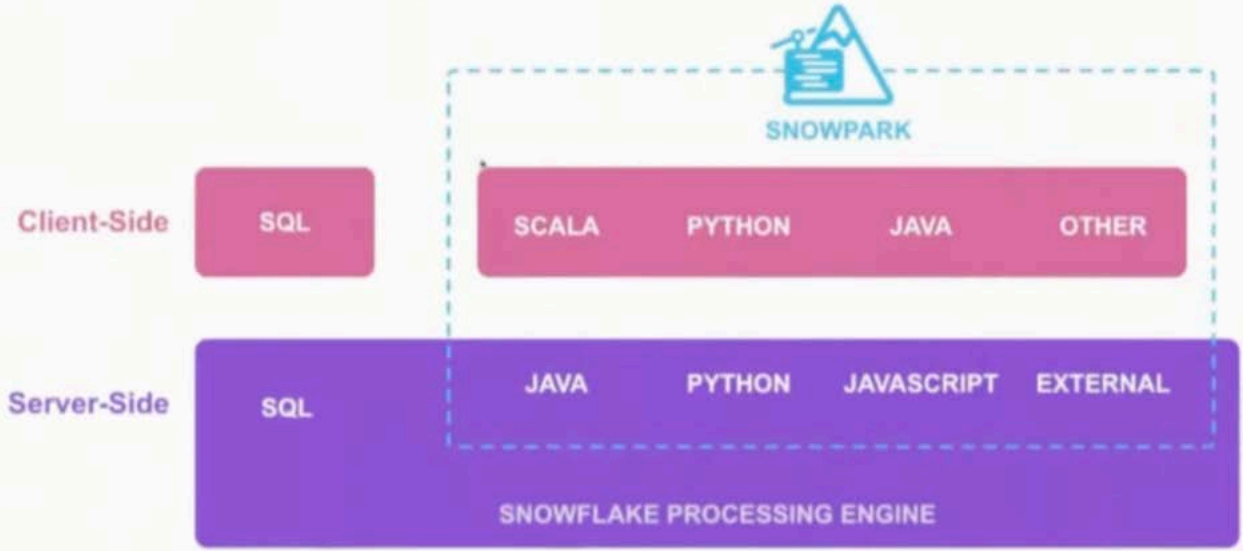

Data Science & Machine Learning Platform

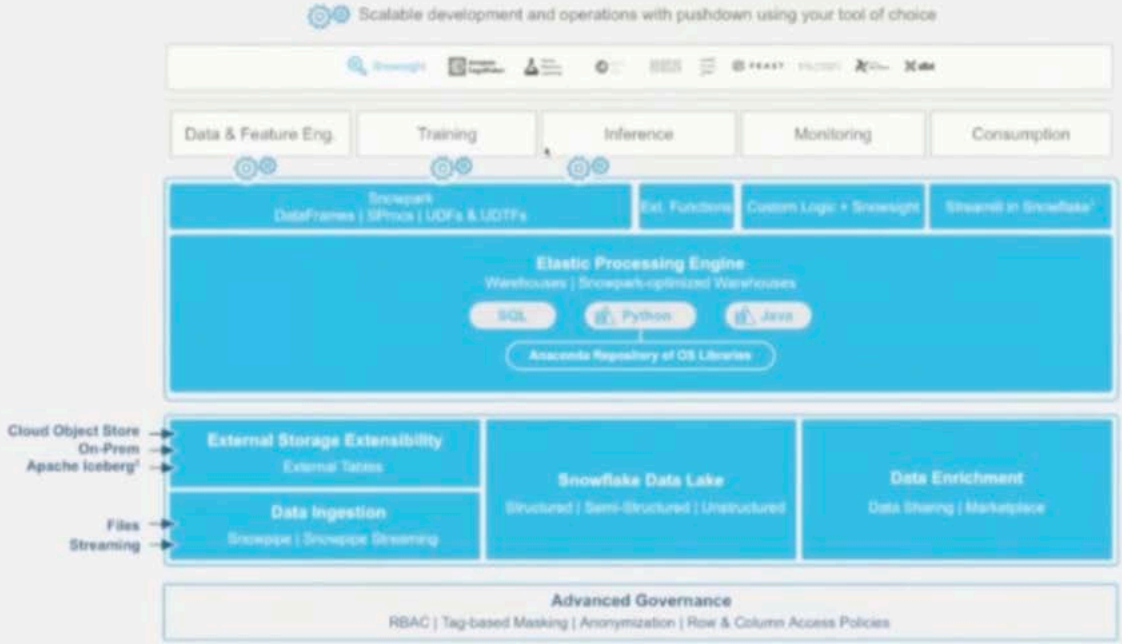

[1] Public preview
[2] Private preview