



Introducing Snowflake

YOUR DATA CLOUD

Regan Murphy, Principal Sales Engineer

regan.murphy@snowflake.com

SNOWFLAKE: YOUR DATA CLOUD

Our vision

Allow our customers to access all their data in one place so they can make actionable decisions anytime, anywhere, with any number of users.



Our solution

Next-generation data platform built from the ground up for the cloud to address today's data and analytics challenges.



A familiar SQL-based platform



Support any analytic workload



Across multiple clouds



Delivered as a service

3 YEARS IN STEALTH + 7 YEARS GA

Founded in 2012 by industry veterans with over 120 database patents

First US customers in 2014
GA in 2015



Over 3500 employees,
6800+ customers today



NZ Office opened
May 2019



Sep 2020
Snowflake (SNOW) lists
on NYSE

FUN FACTS

Queries processed in
Snowflake per day:

> 1,081 million

Largest single
table:

> 48 trillion rows

Largest number of
tables single DB:

2 million!

Single customer
most data:

> 55PB

Single customer most
users:


> 10,000






Why?

IT'S JUST NOT WORKING...




Complexity

Requires management of data, security and infrastructure.



Fragmentation

Multiple disparate platforms create silos of data.




Data Interchange

Sharing requires movement of data which limits volume, frequency and security.



Limited Scalability

Can't support all your data, users and workloads.



Inadequate Elasticity

Rigid, inflexible architectures that cannot flex to match business needs.



Rigid Cost

Based on provisioned capacity models, pay for fixed resource sizing.

How does Snowflake
make things easier?



A NEW DATA PLATFORM ARCHITECTURE

Multi-cluster, shared data, in the cloud

Traditional Architectures



Shared-disk


Shared storage
Single cluster



Shared-nothing

Decentralised, local storage
Single cluster

Snowflake



Multi-cluster, shared data

Separation of compute and storage allows unlimited, independent scalability of each.



Let's dive in...



ETL/ELT



Structured,
Semi-structured



Streaming
Data Lake



Data Pipelines
External Tables



Global Services

- Metadata Repository
- Logical Model
- Security
- Query Planning & Optimisation
- Transactional Control



Reports,
Dashboards,
Analytics



ODBC, JDBC, Python,
Spark, Node.js, Go...





Instant elasticity,
blazing performance,
unlimited scale



Structured,
Semi-structured



ETL/ELT



Data Pipelines
External Tables



Streaming
Data Lake



ODBC, JDBC, Python,
Spark, Node.js, Go...

Reports,
Dashboards,
Analytics



Data
Science




Global Services

- Metadata Repository
- Logical Model
- Security
- Query Planning & Optimisation
- Transactional Control



1. Scale Across - for workload isolation

How do you manage resources across competing workloads today?



Structured,
Semi-structured



ETL/ELT



Data Pipelines
External Tables



Streaming
Data Lake



Global Services

- Metadata Repository
- Logical Model
- Security
- Query Planning & Optimisation
- Transactional Control



ODBC, JDBC, Python,
Spark, Node.js, Go...

Reports,
Dashboards,
Analytics

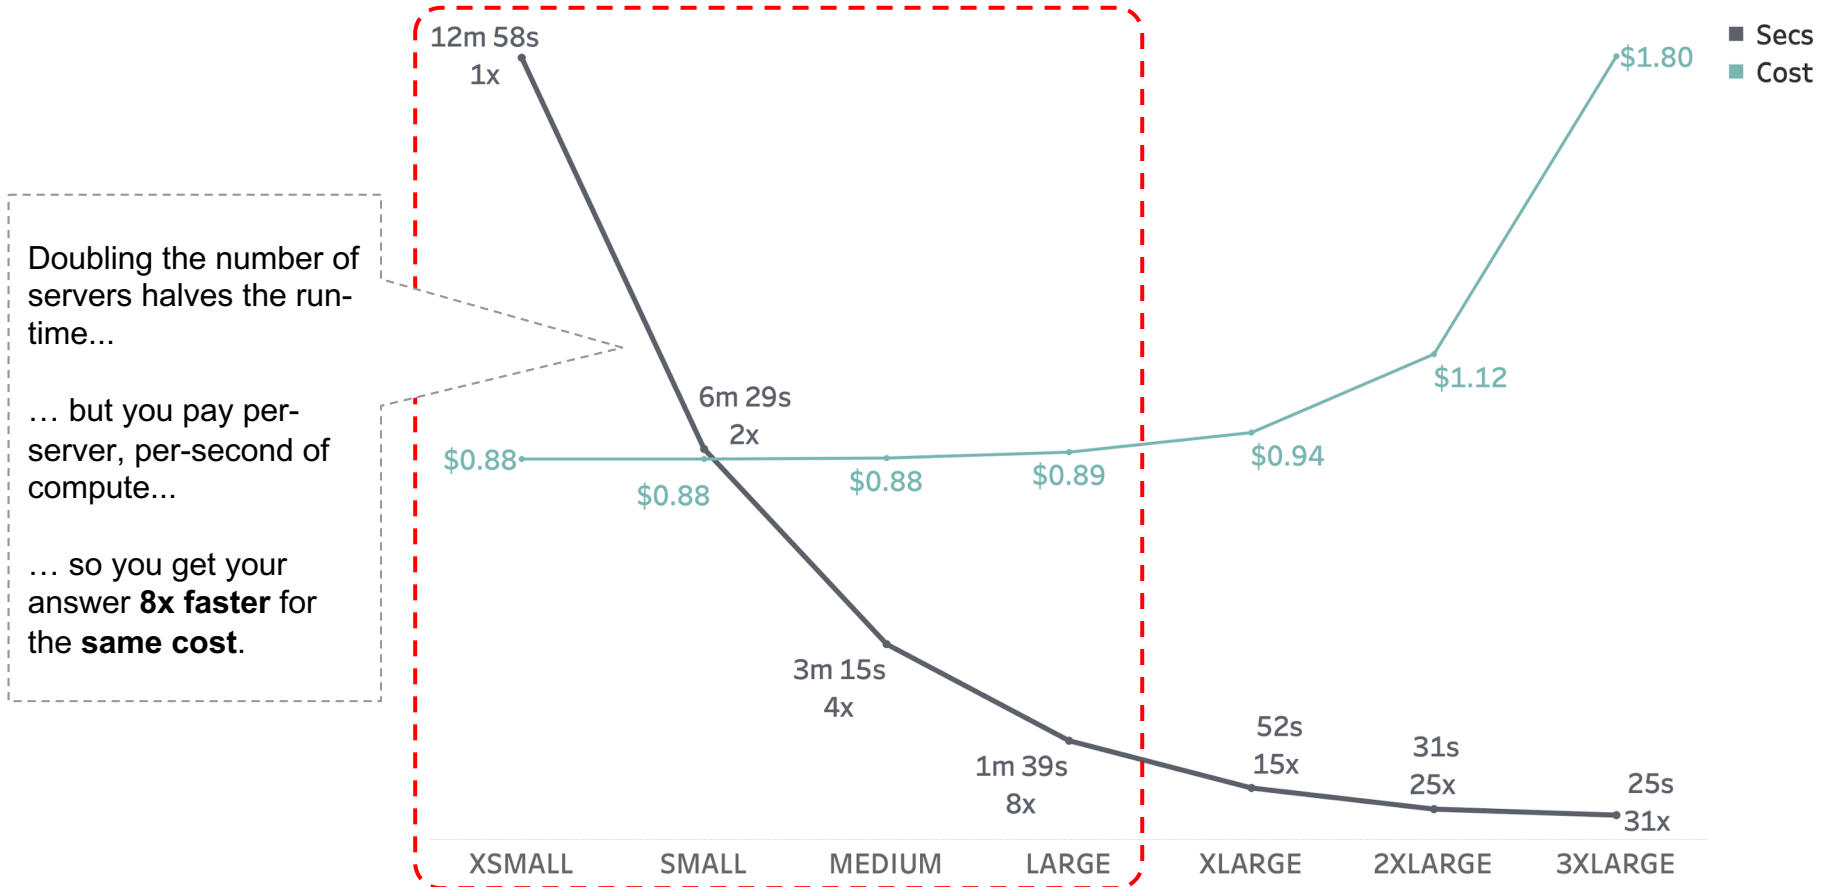



Data Science

2. Scale up - for larger/more complex workloads

How do you handle workloads with varying data volume and query complexity today?

SCALE UP - LOADING 1BN RECORDS





Structured,
Semi-structured



ETL/ELT



Streaming
Data Lake



Data Pipelines
External Tables



ODBC, JDBC, Python,
Spark, Node.js, Go...




Global Services

- Metadata Repository
- Logical Model
- Security
- Query Planning & Optimisation
- Transactional Control



Data
Science




3. Scale Out - for high concurrency

How do you handle spikes/seasonality in user demand today?

PUTTING IT ALL TOGETHER

ETL & Processing



Scale up -
warehouse resize

Data Science



Provision & discard
on-demand

Reporting



Scale out -
multi-cluster
warehouse


Ad-hoc Analytics



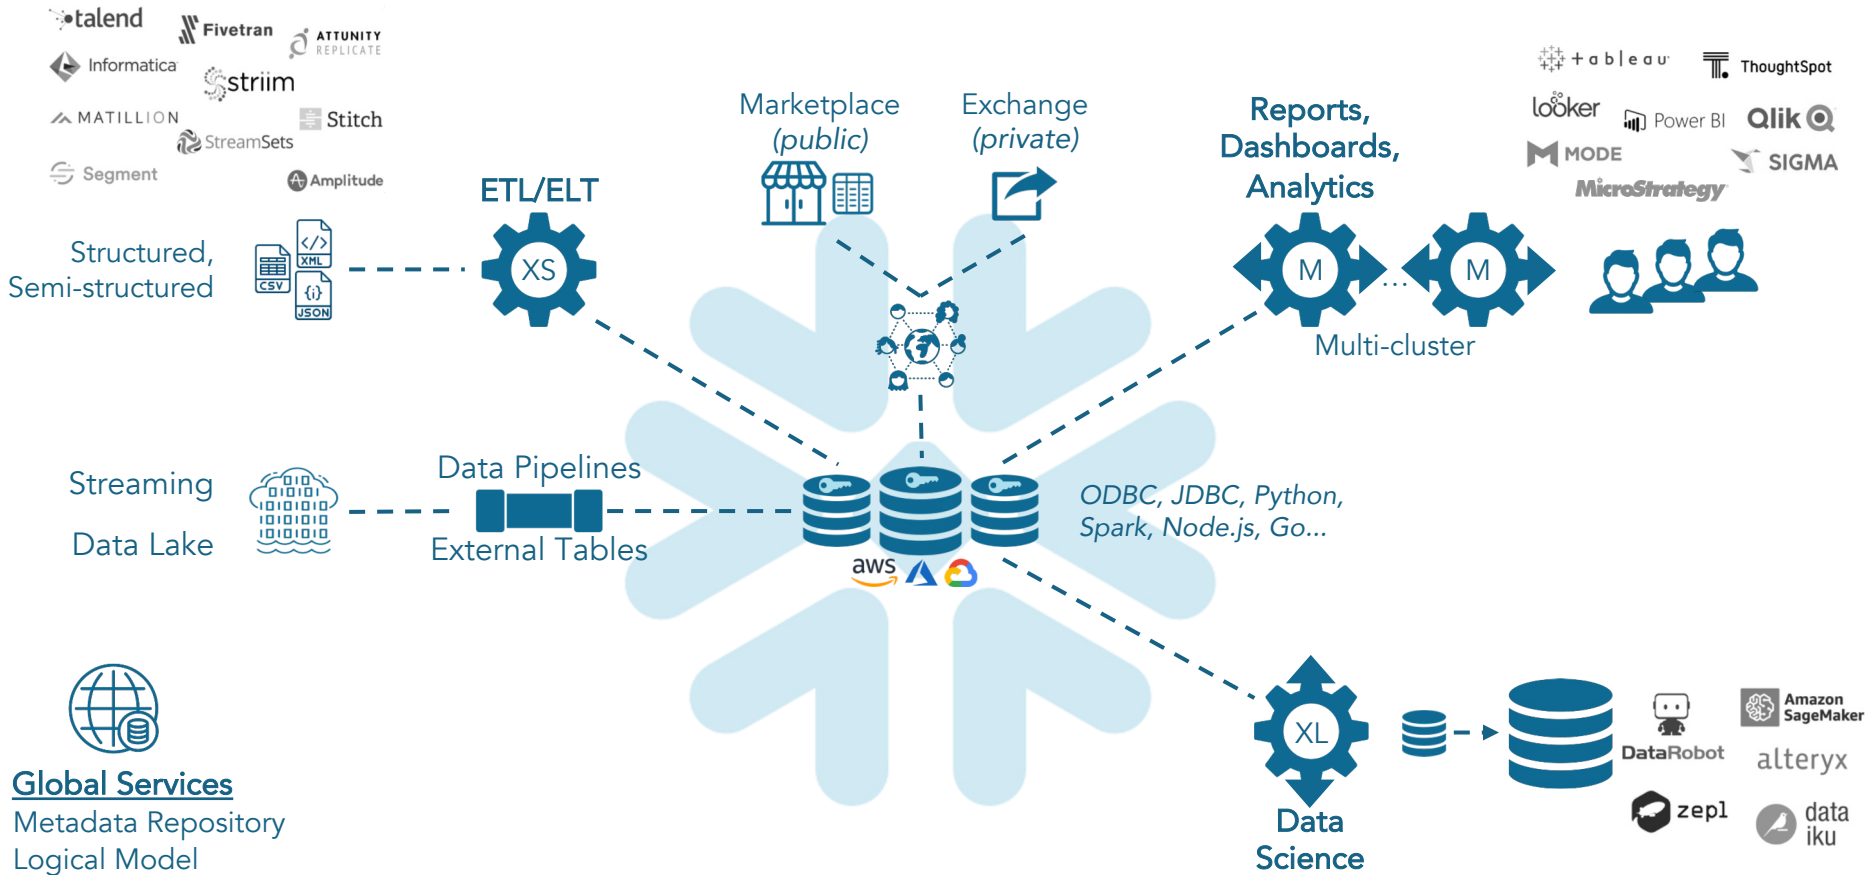
Auto-suspend
& resume

Snowflake handles all of these workloads - at once.

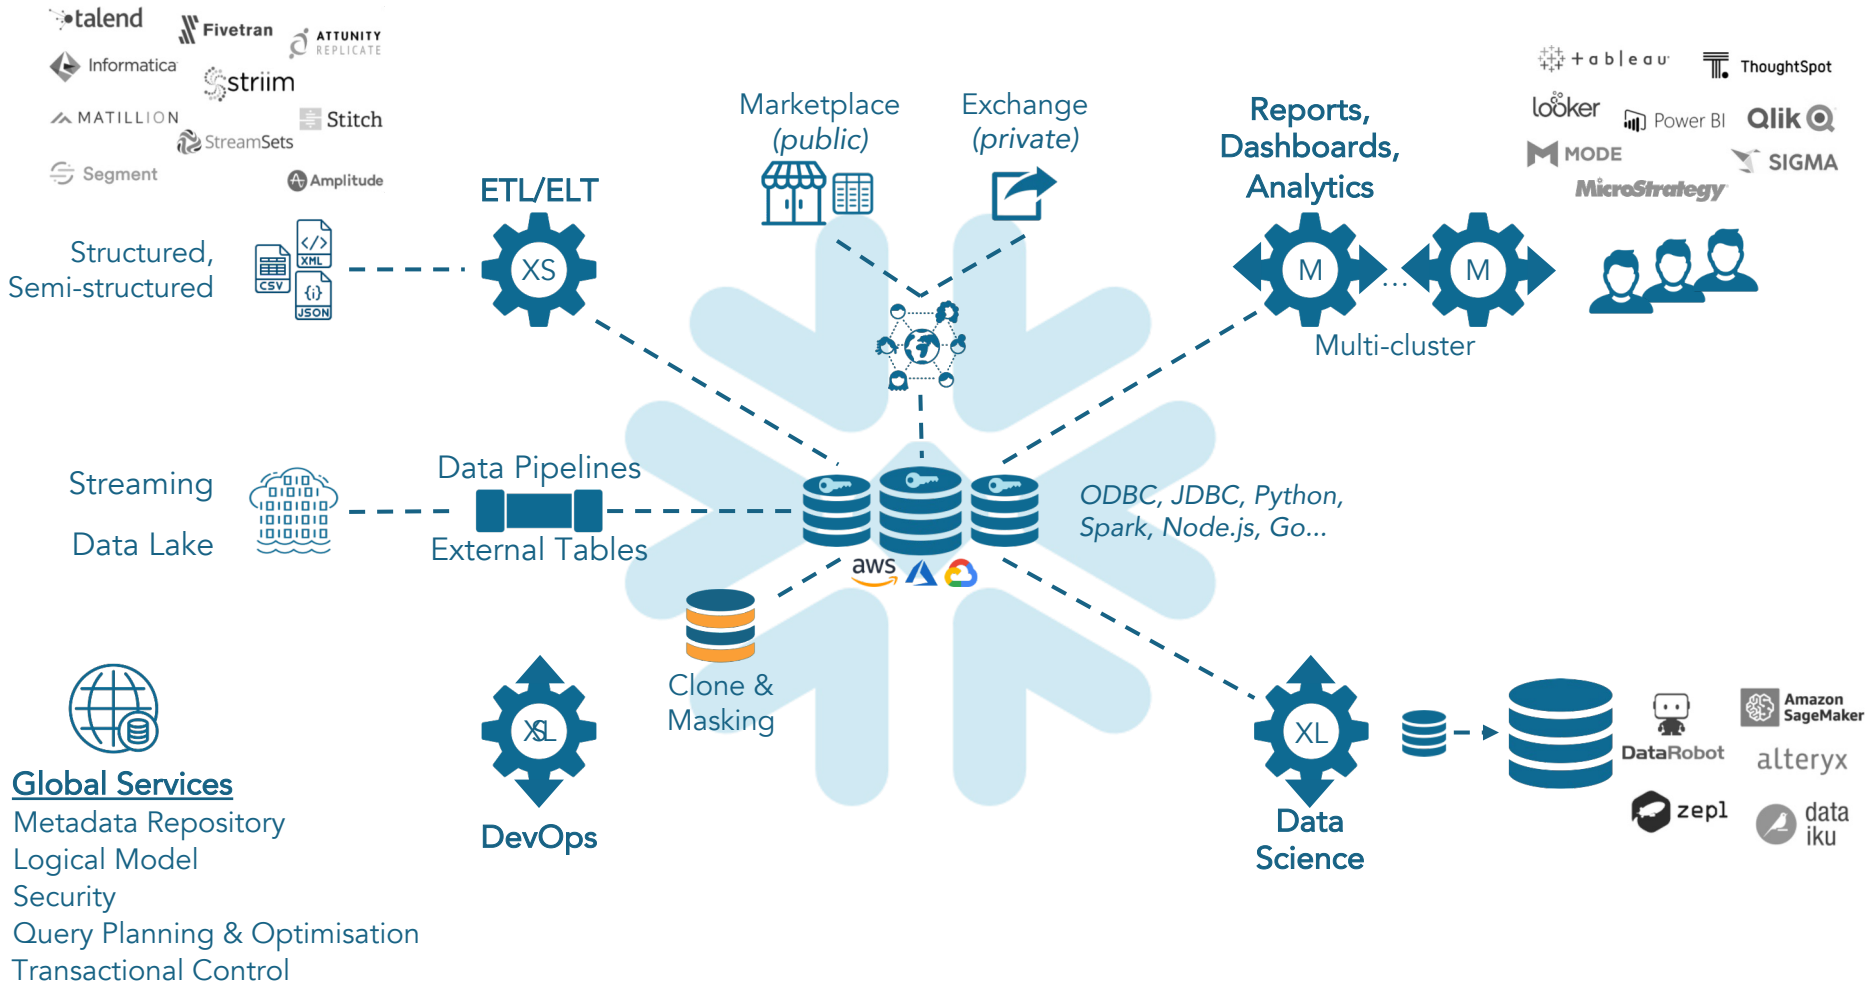
How do you handle multiple, varying workloads today?



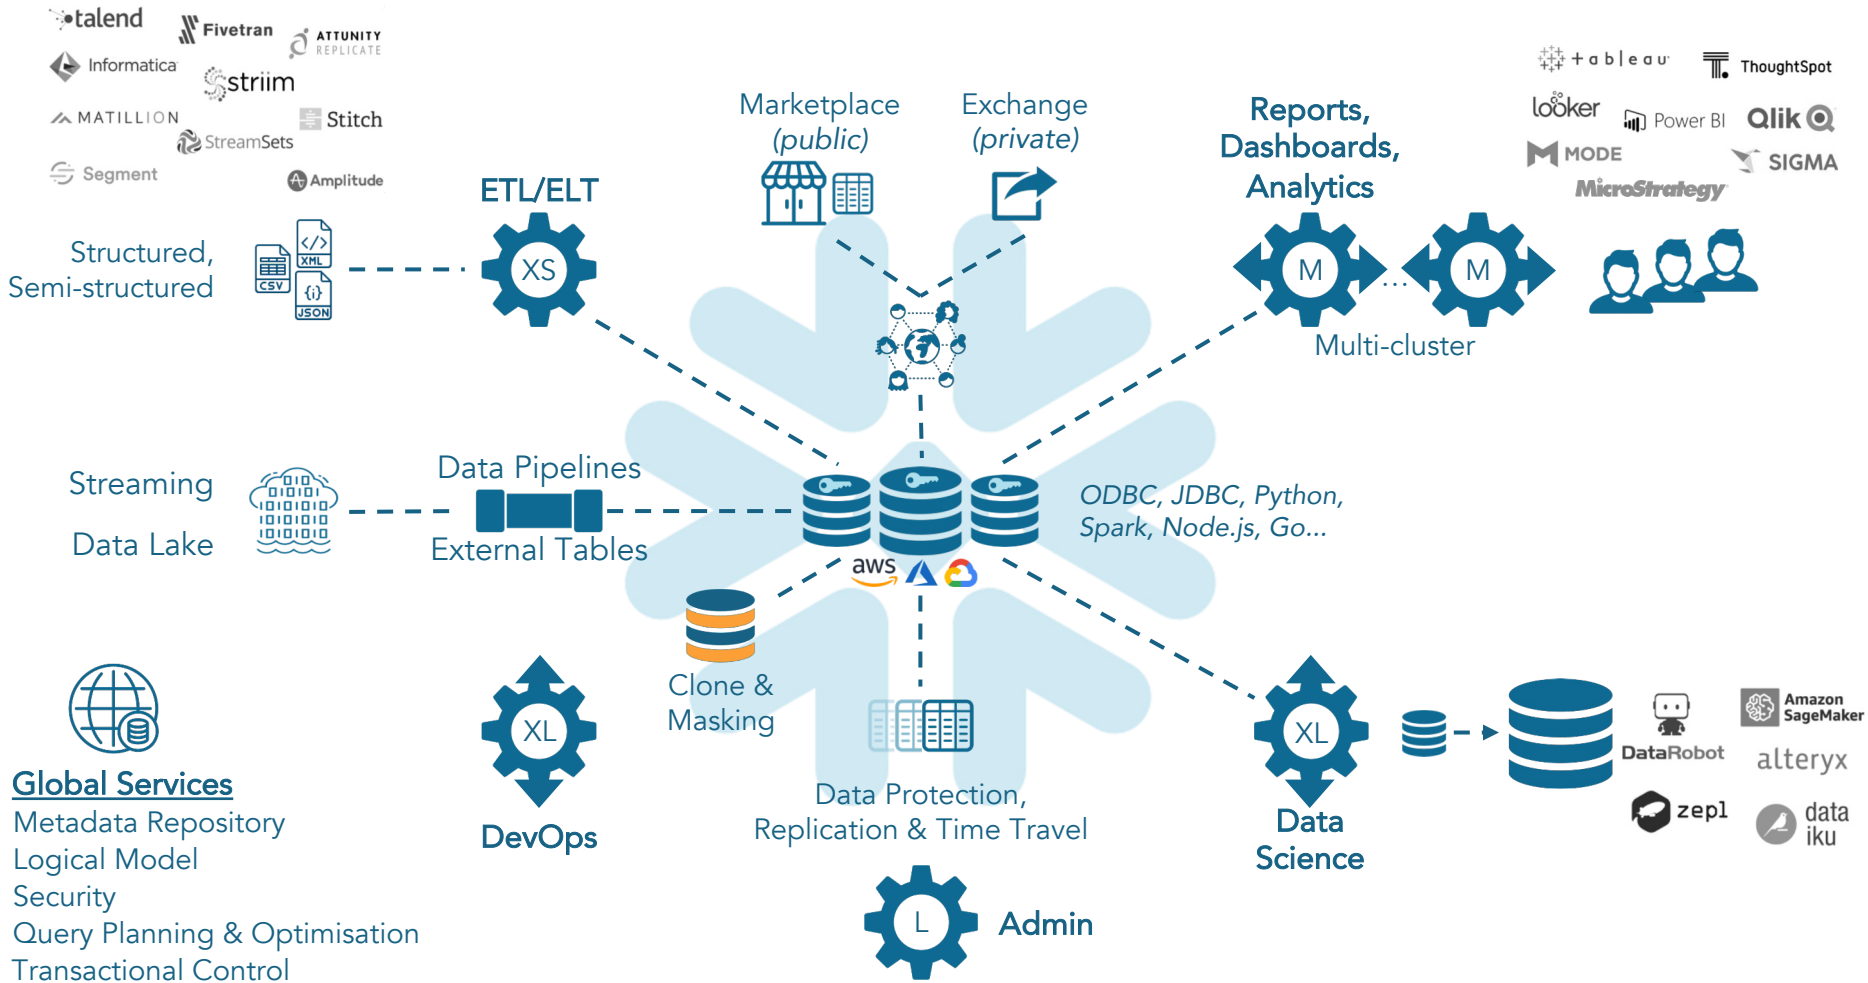
One platform,
one copy of data,
many workloads



How do you handle interchange of data between depts, customers, partners, etc today?



How do you manage non-prod requirements today?




How do you manage backup and restore tasks today?

Wow...
so much to remember...




FROM “IT’S NOT WORKING”...



Complexity

Require management of data, security and infrastructure



Fragmentation

Multiple disparate platforms create silos of data




Data Transfer

Sharing requires movement of data which limits volume, frequency and security



Limited Scalability

Can't support all your data, users and workloads



Inadequate Elasticity


Rigid, inflexible architectures that cannot flex to match business needs



Rigid Cost


Based on provisioned capacity models, pay for fixed resource sizing

TO “IT JUST WORKS”




Simplicity

Simple, serverless, plug and play



Integration

A single platform supporting all analytic workloads




Data Sharing

Secure, live data sharing across accounts, regions and clouds



Unlimited Scalability

Any scale of data, users and workloads. Grow without re-architecting



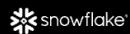
Instant Elasticity

Size for what you need, right now. Resize instantly



Pay As You Go

Pay for what you use, when you use it



WHY SNOWFLAKE

THE DATA CLOUD

SOLUTIONS

RESOURCES

COMPANY



CONTACT

START FOR FREE

JOIN US FOR BUILD ['22] THE DATA CLOUD DEV SUMMIT

Snowpark. Python. Streamlit. Native Apps. Iceberg. Unistore. Explore these topics and more in 30+ sessions packed with demos, labs, and AMAs.

REGISTER NOW



PRICING

WEEKLY DEMO

HANDS-ON LABS

OFFICE HOURS



THANK YOU

