

Snowflake definition:

(n) a flake of snow, especially a feathery ice crystal, typically displaying delicate sixfold symmetry.

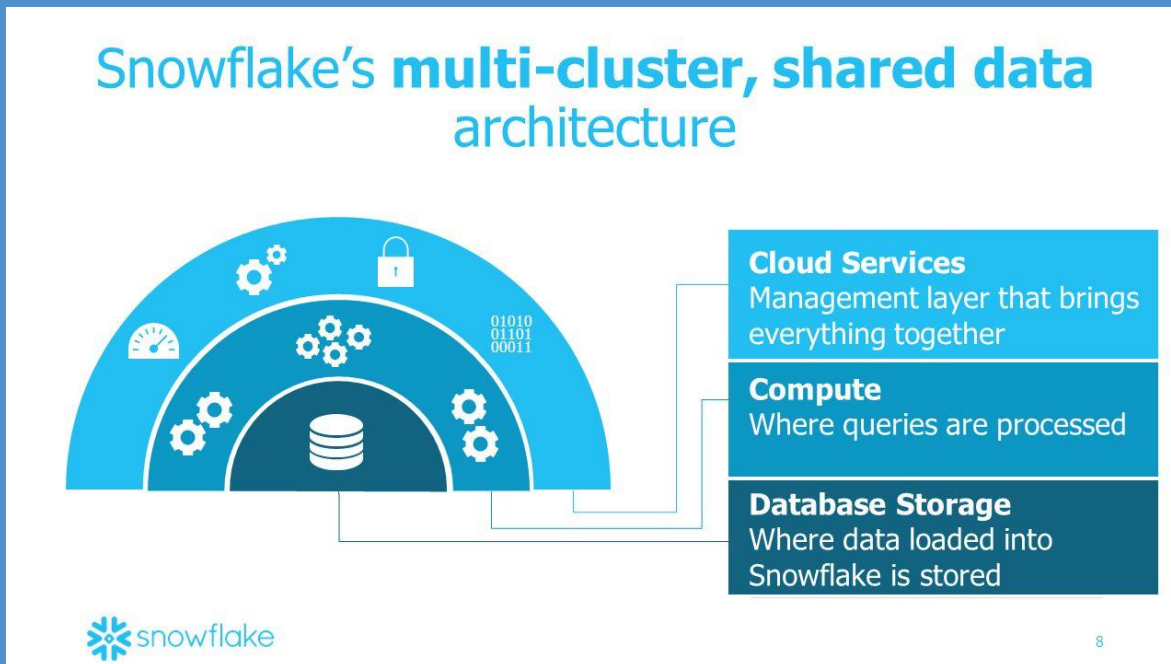


Snowflake definition (according to Tyler Durden of Fight Club):
(adj) A term to describe someone that thinks they are unique and special, but really are not. Gained popularity after the movie "Fight Club" from the quote **“You are not special. You're not a beautiful and unique snowflake. You're the same decaying organic matter as everything else.”**



Snowflake evolved definition (courtesy Snowflake Data Warehouse):

(n) A data warehouse that delivers the performance, simplicity, concurrency and affordability not possible with other data warehouses making users supremely unique and special in their rockstar capabilities!





THE VEHICLE INTELLIGENCE COMPANY

Spireon provides visibility and insights from connected vehicles to help business run smarter, and people drive safer

How Spireon Can Help Your Business



UNIFIED IOT PLATFORM UNDERPINS ALL OF OUR BUSINESSES

~4.5M

ACTIVE SUBSCRIBERS
SUPPORTED

15B+

DATA EVENTS
PER YEAR

99.99%

SYSTEM
AVAILABILITY

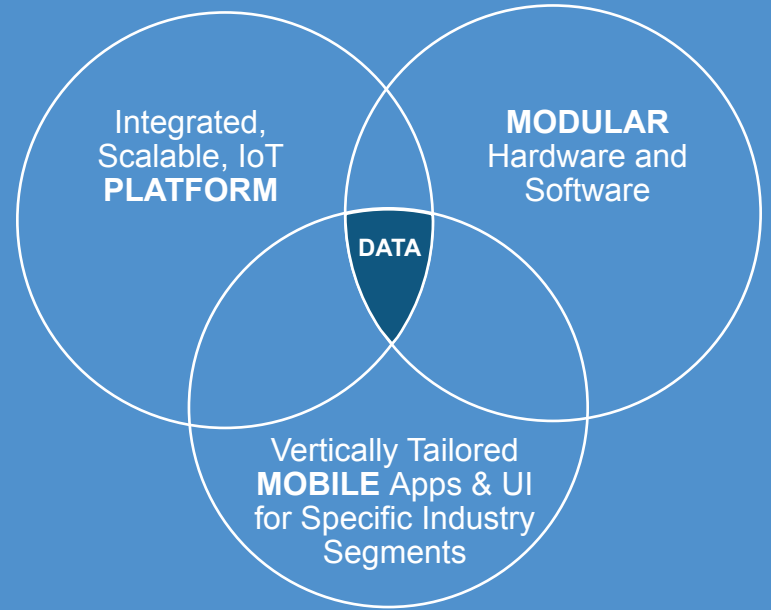
24/7/365

TIER 3 AND 4
DATA CENTERS

Scalable Cloud
Infrastructure

CAN SUPPORT
2B+ USERS

Multi-Layered
Security



CHALLENGES



Database Migration



Data Sharing



Streaming Data

RESULTS

Faster time to insight at reduced cost

- Reduced costly MongoDB footprint
- Broader and real-time data sets
- Insights in < 30 sec vs. 48 hours

Real-time asset tracking for clients

- Enhanced analytics reporting
- Enabled by data sharing
- Differentiated services for clients

New business models and revenue

- Driver scores for usage-based insurance
- New opportunities based on IOT data

Snowflake @ Spireon - Evolution

- Analytics evolution:
 - Reporting from MongoDB for both transactional + analytics
 - Separation of transactional/ODS vs analytics concerns with Snowflake
- Redshift lessons learned
 - Semi-structured data format (e.g. JSON) unfriendly
 - Administrative overhead
 - “Analyze and Vacuum” Hell!
 - Must factor local storage disaster recovery
 - Inconvenient and expensive scaling methods
- ETL evolution:
 - AWS Data Pipeline => Astronomer Airflow => Snowflake Tasks


Snowflake @ Spireon - Today

- Building out Data Warehouse
- Batch
 - Airflow
 - Groovy Scripts
- CDC
 - Alooma
- Schema versions with release process
- Streaming
 - Kafka Connect
 - Snowpipes
- Looker

Kafka Connect & Snowpipe

- Raw Data Payload => Event Archiver
 - Storing payloads sent from devices in the field
 - Used for analytics and debugging for hardware team
 - Used to rely on ElasticSearch with 7 days of data
 - Launched in May 2019
- Performance Stats
 - ~ 14 billion total records in Snowflake
 - ~ 6.4 million records an hour
 - ~ 8.1 min average ingest time Kafka into Snowflake

Kafka Connect & Snowpipe



Snowflake @ Spireon - The Future

- More Streaming
- Use known architectural patterns (like previous slide) across the platform
- Schema Registry
 - Allow for planned and thoughtful schema evolution
 - Keep data lake and warehouse clean
 - Prevent pipelines from breaking
 - Data as a first class citizen
- Self-service data platform
 - Involve full stack developers for complete horizontal implementations
- Evaluating Snowflake as an ODS

Future Opportunity: Snowflake As ODS

- New project: HFTP
- Events API
 - Ingesting 3800 messages / second
 - Reading 30 query / second
 - Append only data
 - Known query pattern
 - Original plan: <6 months in Atlas, >6 months in Snowflake as reports
- Why not just use Snowflake?
 - Is it performant enough?
 - Is it more cost effective?
- Initial results okay, but lots of room for improvements