PERFORMANCE AND SCALABILITY

- Existing ad-hoc querying of the traditional data warehouse are slow in returning results
- On-premise architecture cannot support modern analytics workloads

- Scaling up compute can be performed in an instant with Snowflake
- Clusters can be scaled out automatically and back in based on user demand
- Per second compute billing – no cost for no query activity

Challenges:

As enterprises rise to age of analytics more business users are utilising self-service analytics at their fingertips, where the latest insights are needed at speed of thought.

Current on-premise data warehouses produce bottlenecks, lack optimisation and prioritisation of queries for data driven organisations to get timely insights.

Analytics queries may sometimes take several hours, if not days, competing with other concurrent users and workloads, that utilise the same compute resources.

Solution:

There are various ways to tackle the query performance problems across the existing on-premise data warehouse and data lakes on an organisation's existing platform.

One use case would be to provision more hardware for compute and storage for the existing platforms.

Provisioning more hardware is a costly, resource-intensive activity that may well be repeated in the near future, as analytics workloads and the userbase grows exponentially.

Moreover, the unpredictability of future data volume, variety and velocity of workloads (e.g. machine learning and AI) require a total re-architecture and provisioning compute and storage resourcing.

Traditional on-premise architectures don't have the flexibility nor scalability to deal with the demand of a modern data architecture. A new architecture is needed to support the flexibility of instant scale up and out during peak workload activity

In the age of the cloud where storage and compute have become relatively affordable, a cloud data platform allows for a cost-effective controlled and almost instant way to provision demanding cross-functional workloads.

How Snowflake can help?

Snowflake empowers all users to leverage performance as required. Workloads for analysts and data scientists can be scaled up on demand as data volumes and the complexity of queries increase and can scale back down during normal business operations.

Snowflake's unique architecture with the separation of compute and storage allows for instant allocation of processing power to specific use cases and workloads, while billing to the second.

Snowflake has the ability to instantaneously increase/decrease its clustered workload on-demand for each virtual warehouse, whilst data is independent, a single copy can be shared to various parts of the organisation and even external parties.

Snowflakes 'Virtual Warehouses' can be provisioned in an instant.

In use cases where there are multiple concurrent users during certain hours, e.g. marketing or sales end of month reporting, Snowflake's automated multi-cluster configuration can deal with the breadth of multiple concurrent users, accessing the same data sets, then automatically scaling down the multiple clusters once user queries have subsided.