

REDUCED MAINTAINENCE AND DATA PROTECTION



- Near zero management – no knobs to turn pure 'software as a service' data platform
- No indexes, partitions, distribution keys required – just simple load the data to your table
- Scale on demand – scale your compute, storage and user concurrency in an instant

Challenges:

As volumes and variety of data ever increases; so too does the complexity of maintaining data on traditional data platforms.

Cumbersome maintenance activities, with a vast array of resources, are required for frameworks such as HADOOP, which makes for an overly complex and costly data environment.

Resources required in ageing data platforms include data engineers, full stack programmers, solution architects and platform engineers across multiple ecosystems, in most cases with a mix across on-premise and cloud infrastructure.

Technical complexities include partitions on your Hadoop file system and file maintenance. Moreover, data engineers are constantly challenged with indexing, partitioning and collecting statistics for tables within your data warehouse.

Legacy on-premise data platforms require outdated maintenance activities, where backup to file or tape are cumbersome for data restoration and may take days or even weeks. These factors effect downtime of a data platform, which have become mission-critical in modern data driven business.

Solution:

Data platforms have evolved over the last few years, namely within the cloud where vendors have been able to re-architect the way data is ingested, consumed and maintained on a single platform without the burdensome maintenance, disaster recovery and failover activities. We now see capabilities such as cross-region replication and recovery that cannot be rivalled with on-premise outdated hardware-technology.

Some of the leading next generation cloud data warehouses provide automation on indexing/pruning, automated backup and recovery activities as well resolving



- Restore snapshot of database objects in an instant with Snowflakes unique 'time travel' and 'cloning' features
- Disaster recovery and 99% uptime can be reassured with Snowflake's cloud vendor partnerships
- Database replication and failover can be achieved via multiple Snowflake accounts setup across regions and different cloud providers.

user generated data corruption, at the click of a button without platform administrators involvement.

How Snowflake can help?

Snowflake's unique technological features, such as micro-partitioning for self-adopting query performance and distribution, eliminates the need for data engineers and database admins to plan and execute complex indexing, partitioning, distribution and skew on tables.

This is all taken care of by Snowflake, as the data is naturally ingested regardless of whether it is structured or semi-structured data.

Snowflakes 'continuous data protection' features allows user and system driven data corruption to be restored in an instant, with a combination of Snowflake's 'Time Travel' and 'Cloning' features.

Continuous Data Protection Lifecycle



Time Travel allows users to retain historical data that may have been updated or deleted at any point in time, for 90 days

Cloning creates metadata, copies database, schemas and tables to create a snapshot at any point in time, that can be later historically queried with simple SQL predicates.

In addition to continuous data protection, Snowflakes cloud partners offer regional failover to guarantee no downtime.

In addition, Snowflake offers enterprise-grade database replication, that can be performed across different Snowflake cloud providers, accounts and regions.