snowflake®

Snowflake High Availability for Data Applications

Introduction

This white paper is intended for anyone who is or will be developing a data application on Snowflake. The first section describes how Snowflake defines and calculates availability along with the service level that can be achieved out of the box. The remaining sections focus on Snowflake features that can be utilized for applications that demand higher availability service levels.

Availability in the Cloud

Snowflake was built in the cloud, for the cloud, to take advantage of all the features the cloud has to offer.

Cloud infrastructure

Cloud infrastructure makes possible new capabilities for ensuring availability. It provides easy access to multiple data centers connected by high-performance networks, where each data center contains a huge pool of low-cost, commodity hardware. Cloud infrastructure also offers scalable storage services designed for high availability and durability, on-demand virtualized compute, and high-speed network connectivity to connect them (commonly 10GigE network connections with even higher-capacity backbones).

The low cost and on-demand availability of infrastructure in the cloud makes it practical to design software where every component leverages multiple compute instances, providing increased resiliency to server failures. Cloud storage services also make it easy to take advantage of data redundancy and replication, even across multiple data centers.

To help you fully understand Snowflake's built-in resilience and fault tolerance, let's first review some terminology.

Availability: This metric measures the uptime of a service or the time during which service operations run without a service unavailable error. This metric is usually calculated as a percentage such as 99.9% on a monthly basis based on the ratio of uptime minutes in the month over total minutes in that month. Table 1 shows the amount of time the resource can be down before it breaks the availability service level agreement (SLA). Make sure to

review and evaluate your cloud provider SLA calculation to understand what they commit to.

Availability Service Level	Maximum Downtime per Year
99%	3 days 15 hours
99.9%	8 hours 45 minutes
99.99%	52 minutes
99.999%	5 minutes

Table 1: Availability service level

Durability: This metric defines the probability that there is no data loss. It is usually defined as a percentage such as 99.999999999% over the course of a year.

Snowflake editions (as of May 2021)

Snowflake offers multiple editions, ensuring that your usage fits your organization's specific requirements. Each successive edition feature set builds on the previous edition.

Standard Edition: Offers introductory-level capabilities providing full, unlimited access to all of Snowflake's standard features. It provides a strong balance between features, level of support, and cost.

Enterprise Edition: Offers features that enable organizations to manage multiple accounts, giving them the ability to have multi-cluster warehouses, materialized views, and column- and row-level security.

Business Critical Edition: Offers even higher levels of data protection to support the needs of organizations with extremely sensitive data, particularly PHI data that must comply with HIPAA and HITRUST CSF regulations. Database failover/failback adds support for business continuity and disaster recovery.

Virtual Private Snowflake (VPS): Offers our highest level of security for organizations that have the strictest requirements, such as financial institutions and any other large enterprises that collect, analyze, and share highly sensitive data.

Snowflake Architecture for Availability

Snowflake was designed from the start to be a highly available data platform in the cloud and achieve an availability service level of 99.9% or greater across all editions. It has built comprehensive fault tolerance and resiliency into its product, helping customers ensure that the data and its data platform are always available while removing some of their infrastructure management responsibilities.

That design begins with the Snowflake architecture. The Snowflake architecture consists of the following three layers, each decoupled from the others and distributed across multiple data centers to provide failure resiliency.

- **Centralized Storage**: This layer uses cloud object storage from AWS, Microsoft Azure, and GCP to persist all database data.
- **Multi-Cluster Compute**: This layer contains multiple massively parallel processing (MPP) compute clusters that process queries.
- **Cloud Services**: The overall brain in the system, this layer is a collection of services that handle query management, optimization, transactions, security and governance, metadata, and sharing and collaboration.

All of these layers run within a region of AWS, Microsoft Azure, or GCP. Each region consists of multiple availability zones. An availability zone is one or more physically separated data centers with independent access to power and networking. The Snowflake deployment extends across three availability zones. The Snowflake storage layer is synchronously and transparently replicated across multiple disk devices and availability zones by the cloud object stores. In addition, Snowflake's virtual warehouses and cloud services layers are similarly deployed across three availability zones in a region.

Data protection and data availability

The database storage layer of Snowflake is the encrypted repository for all customer data. Snowflake stores all persisted data in the cloud object store of the cloud provider for your deployment. Transactions are committed only after all modified data has been written to cloud storage. Cloud storage in turn synchronously and automatically replicates the stored data across multiple devices and across at least three availability zones. Snowflake is built on the cloud providers' blob stores which all tout at least 11 9's of durability.

Virtual warehouses and query resiliency

Snowflake's architecture also ensures that queries complete successfully in the face of a broad array of failures. Query execution tasks are processed in the virtual warehouses. If a single compute instance fails within a virtual warehouse, Snowflake automatically replaces that instance and completes the query without disruption. Virtual warehouses do not have permanent storage of data. Instead, they cache data to improve query performance. Loss of any number of nodes in a virtual warehouse, including the loss of an entire virtual warehouse, does not cause a data loss since the persistent copy of data resides in the cloud object stores.

For performance reasons, each virtual warehouse runs within a single availability zone. However, Snowflake can provision a virtual warehouse within any available availability zone within a region. If Snowflake detects that an availability zone has become unavailable, the cloud services layer can reprovision the impacted virtual warehouse in a different zone and restart queries that have been in process to ensure that they complete successfully.

Service availability

Snowflake's cloud services layer is responsible for handling requests from connected clients and for providing services including security, virtual warehouse management, query optimization, and metadata management. This layer is also designed for resiliency, ensuring that the Snowflake Data Cloud is fault tolerant across a comprehensive set of failure scenarios.

The cloud services layer takes into account the physical availability zones of Cloud Service Providers. The Cloud Services Layer is designed to actively consider this in its scheduling/placement algorithm to enforce redundancy across zones.

The services in the cloud services layer rely on a metadata storage system that keeps track of all metadata needed by Snowflake, from information about where database data is stored to information about virtual warehouse configurations and query execution history. This metadata storage system is implemented to be both transactionally consistent and ACID compliant. It is deployed across multiple virtual compute instances such that all updates are synchronously committed to multiple instances in multiple availability zones. In the event of a failure of one or more instances supporting this

metadata storage system, or even the loss of up to two data centers, the metadata storage system can continue operation without disruption and without loss of data.

Near-zero maintenance

Snowflake provides for continuous operation not only when failures occur, but also when Snowflake performs software updates. This means that customers do not have to spend time or effort to maintain Snowflake unlike other infrastructure-hosted solutions which require downtime for traditional database administration activities.

Going Beyond 99.9% Availability

Customers may require greater than 99.9% availability from Snowflake for their applications. Snowflake provides additional features that can enable customers to achieve higher availability.

Database replication and failover

The features discussed so far—multi-zone deployment, query resiliency, and service availability—enable your Snowflake account to achieve 99.9% availability.

Customers who want to achieve higher availability can leverage a number of Snowflake capabilities to do so. For example, if a cloud region outage affects Snowflake, customers can replicate their data to another Snowflake account in a different region. In addition, customers can replicate to another cloud provider including GCP, Microsoft Azure, and AWS. The replicated databases are read-only.

Read-only replicas are supported on Snowflake Enterprise Edition or higher.

If your application requires a read-write enabled database, then you can initiate failover by promoting a secondary replica database in an available region to serve as the primary database. When promoted, the now-primary database becomes writeable. Concurrently, the former primary database becomes a secondary, read-only database.

Database failover requires Snowflake Business Critical Edition or higher.

Client redirect

In addition to initiating a failover for the database, Snowflake provides a feature called *client redirect*. This shortens the recovery time for client applications during an outage from potentially hours to mere seconds. The process is quite simple and requires a connection URL to be created for Business Critical accounts using the Organizations feature. Client redirect allows for clients to be redirected automatically within seconds of a failover occurring during an outage.

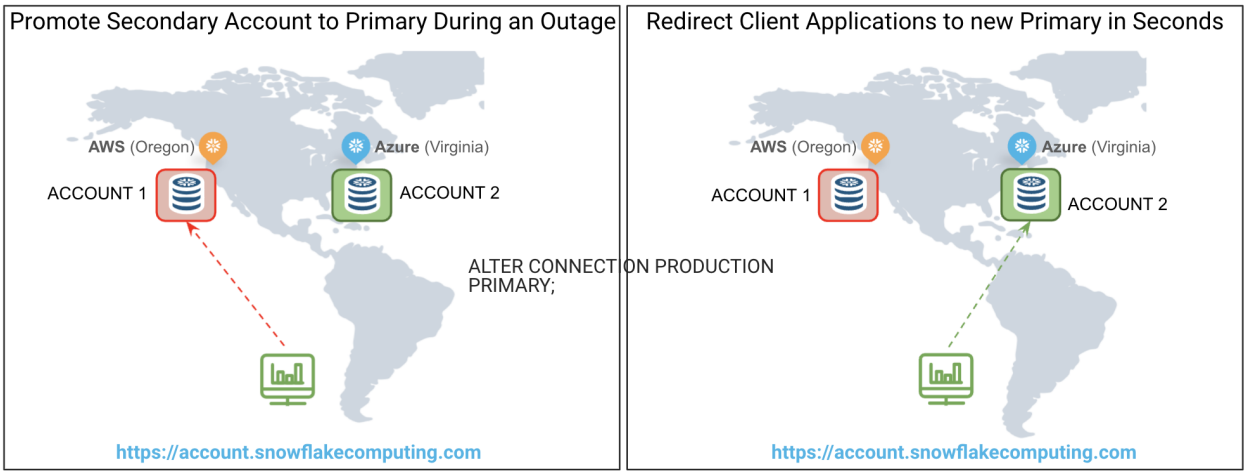

Figure 1: Client redirect

Conclusion

Snowflake provides a level of resiliency and availability that would be so complex and expensive to implement with traditional solutions that it would be impractical for most organizations to achieve. Snowflake leverages the cloud to ensure redundancy and transparent failure recovery, whether from failure of a single component or even of multiple availability zones. This resiliency is built into the Snowflake service and is available out of the box without setup or management. That brings enterprise-class 99.9% availability and disaster recovery for all customers using Snowflake. In addition, by taking advantage of database replication, database failover and client redirect, you can attain higher availability from your application.