

# Extend Snowflake Data Access Securely and Cost Effectively



## Key Benefits

- Manage Snowflake data records on-premises, at the edge, or in the cloud.
- Meet SLAs and comply with regulatory mandates for data privacy and retention.
- Massively scalable platform for data consolidation and easy management.
- Seamless S3-based integration of Snowflake with Cohesity SmartFiles.
- Software defined hybrid cloud flexibility with limitless scale.
- Lower TCO with data consolidation and dramatic data reduction up to 96x or more.
- Robust cyber resiliency and ransomware protection to defend data and applications.

## The Situation

In recent decades, enterprises have relied on legacy data warehouse providers to manage volumes of data for analytics in the data center. Such solutions have been enhanced to provide cloud benefits but remain challenged by architectural limitations associated with their data center focused heritage. With the rise of cloud-based services, businesses are increasingly considering modern cloud-native offerings providing accelerated query performance, independent optimization of compute and storage, flexible pricing options, built-in AI/ML, multicloud support, and more. However, organizations are each pursuing different journeys towards cloud adoption, and not all data is appropriate for management in the cloud. For governance, compliance, and data access performance reasons, many organizations must also maintain their sensitive data on-premises, which challenges the ability to analyze data holistically and efficiently wherever it resides.

## A Modern Cloud-Native Solution

Snowflake was developed to be cloud-native from the start. The Snowflake Data Cloud unites siloed data, discovers and securely shares data, and executes diverse analytic workloads, while delivering a single data experience that spans multiple clouds and geographies. In doing so, Snowflake breaks down data silos and helps derive value from rapidly growing data sets in secure, governed, and compliant ways.

## Why Snowflake and Cohesity for Extending Your Snowflake Data Cloud?

For data that cannot live in the cloud due to data sovereignty requirements, file size, or other limitations, Snowflake supports external tables where data can be stored in an external stage located on-premises. External tables store file-level metadata about the data files, which enables data stored in external stages to be queried as if it was managed by Snowflake in the cloud. Together, Snowflake and Cohesity SmartFiles deliver a joint solution to allow Snowflake data lakes to access data records in the cloud, at the edge, and within on-premises locations.

SmartFiles provide a modern, massively scalable platform for managing unstructured data, with a singular interface that simplifies data management, resiliency, privacy, and control. It is a software-defined, data-centric, unified file and object solution for the hybrid cloud with native S3, NFS, and SMB protocol support. SmartFiles simplifies long-term unstructured data retention, and provides cost-effective, secure, and immutable storage for your Snowflake datasets. Organizations are uniquely empowered to efficiently manage data at scale, with dramatic space efficiency for significant cost of ownership savings. For optimal data placement in the cloud, SmartFiles' built-in hybrid cloud flexibility provides integration with leading public cloud services. Microsoft Azure, Amazon AWS, and Google Cloud Platform can be used for long-term data archival, dev/test, and disaster recovery – allowing you to take advantage of the scalability and pay-as-you-go model of the cloud.

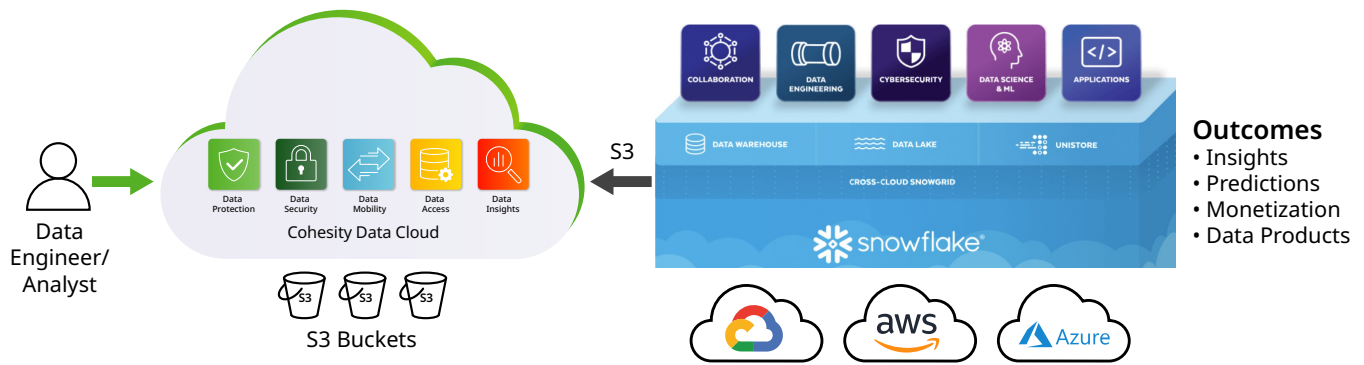


Figure 1: Securely augment your cloud data lake with Cohesity and Snowflake

## Joint Solution Features and Capabilities

- **Enterprise cloud data lake:** Snowflake’s cross-cloud platform breaks down silos by supporting a variety of data types and storage patterns. Users across organizations can securely access structured, semi-structured, and unstructured data for a variety of workloads, with the performance they expect from an enterprise data solution. Data is stored in Snowflake with efficient compression, automatic micro-partitioning, and encryption at rest and in transit. Data can also be accessed from SmartFiles, which provides multiple benefits:
  - Query on-premises data lakes without ingesting data into Snowflake.
  - Create materialized views on external tables to speed query performance.
  - Run ad-hoc queries directly on raw data before ingesting the data into Snowflake.
- **Seamless integration:** Snowflake integrates with SmartFiles using the S3 API to provide persistent, simple, reliable management of data in Snowflake external tables.
- **Lower costs for protection and recovery:** When application data is consolidated on SmartFiles, you can reduce storage requirements by as much as 96x or more through data deduplication, compression, and small file optimization across data center volumes. Storage costs and performance can be further optimized by intelligently moving data to the most appropriate locations or clouds.
- **Robust security and compliance with threat protection:** SmartFiles uses Zero-Trust principles to provide an easy to manage, multi-layered, enterprise security approach to prevent, detect, and analyze threats to your data, including ransomware.

Limit unnecessary data access and prevent unauthorized data addition, modification, and deletion with role-based access controls. Prevent threats and maintain compliance with an immutable file system, encryption at-rest and in-flight, multi-factor (MFA) and quorum authentication, and DataLock (WORM). Integrated anti-virus detects infected files to prevent files from being compromised. Built-in analytics features allow you to observe user behavior analytics for security and compliance investigations. Cohesity FortKnox provides an additional safeguard against ransomware by virtually air-gapping your SmartFiles data to a Cohesity-managed cloud vault for immutable protection with rapid recovery.

- **Hybrid cloud flexibility:** Achieve massive, seamless scale for your SmartFiles data with software-defined simplicity and broad flexibility to deploy on-premises, at the edge, or in your cloud of choice. Manage all your data with a singular management pane with policies to tier, archive, and replicate data across one or more clouds.

## Optimize Access, Cost and Compliance for Snowflake Data Management

Cohesity SmartFiles is the smartest way to optimize cost, scale, and efficiency for your Snowflake data.

Easily manage, secure, and do more with your Snowflake data with a validated solution featuring intelligent data mobility. Realize the benefits delivered by Snowflake with a hybrid cloud that lets you be more resilient, compliant, and cost efficient. SmartFiles’ dramatic scale, flexibility, and operational efficiencies also make managing, protecting, and recovering your data easier than ever. And with a powerful, multilayered security approach, your data can be protected from sophisticated cyberthreats now and into the future.

Make your vision for your Snowflake Data Cloud a reality with **Cohesity**

**COHESITY**



© 2023 Cohesity, Inc. All rights reserved.

Cohesity, the Cohesity Logo, SnapTree, SpanFS, DataPlatform, DataProtect, Helios, the Helios logo, DataGovern, SiteContinuity, and other Cohesity marks are trademarks or registered trademarks of Cohesity, Inc. in the US and/or internationally. Other company and product names may be trademarks of the respective companies with which they are associated. This material (a) is intended to provide you information about Cohesity and our business and products; (b) was believed to be true and accurate at the time it was written, but is subject to change without notice; and (c) is provided on an “AS IS” basis. Cohesity disclaims all express or implied conditions, representations, warranties of any kind.