

Secure Data Sharing in Financial Services: Redshift Data Sharing vs. Snowflake Secure Data Sharing

Ujjawal Nayak¹, Surbhi Gupta²

¹Software Development Manager, Costa Mesa, CA

²Specialist Credit Risk, Seattle WA

Abstract

The **financial services industry** operates in a highly regulated environment that demands **secure, compliant, and efficient data exchange** to enable critical functions such as **credit scoring, fraud detection, risk analytics, and regulatory reporting**. As institutions accelerate their **cloud transformation journeys**, the ability to **share data seamlessly across business units, partners, and ecosystems**—while maintaining strict control and auditability—has become a strategic imperative.

Modern **cloud data platforms** such as **AWS Redshift** and **Snowflake** have introduced advanced, **native data-sharing capabilities** that eliminate the need for data duplication and complex ETL pipelines. These features enable **real-time collaboration** across teams and external stakeholders without compromising on **security, governance, or compliance**.

This article provides a **comprehensive comparison** between **Amazon Redshift Data Sharing** and **Snowflake Secure Data Sharing**, analyzing their capabilities in **security and encryption standards, data governance and lineage, performance and scalability, and regulatory compliance** (e.g., **GDPR, CCPA, GLBA, SOX**). The discussion highlights architectural differences, operational trade-offs, and alignment with financial-sector compliance frameworks.

I. Introduction

Financial organizations operate under strict regulations such as **GDPR, CCPA, SOC 2, and GLBA**, requiring secure handling of sensitive data like credit histories, transactions, and personally identifiable information (PII). Traditional methods of data sharing—such as **ETL pipelines or file-based exchanges**—are inefficient and prone to compliance risks. Cloud platforms now offer native mechanisms for secure, near real-time data collaboration without data duplication. Two prominent solutions are **AWS Redshift Data Sharing** and **Snowflake Secure Data Sharing**.

II. Redshift Data Sharing

Redshift Data Sharing allows multiple Redshift clusters (producer and consumer) to access the same data without duplication.

A. Architecture

- Data is shared at the schema, table, or view level between Redshift clusters.
- Consumers query live data, eliminating ETL or file transfers.
- Designed for AWS ecosystem integration (Glue, S3, Lake Formation).

B. Security and Governance

- IAM roles and AWS Lake Formation policies enforce fine-grained access.
- Encryption at rest/in transit via AWS KMS.
- Resource-level sharing restricted to clusters within the same AWS account or organization.

C. Performance and Cost

- Eliminates data duplication but may introduce inter-cluster network overhead.
- Reserved Instances and concurrency scaling mitigate cost issues.

D. Limitations

- Limited cross-region and cross-account sharing.
- Governance depends heavily on AWS account structures and policies.
- Not inherently multi-cloud.

III. Snowflake Secure Data Sharing

Snowflake Secure Data Sharing, part of its Snowgrid architecture, enables data exchange across accounts, regions, and even cloud providers.

A. Architecture

- Zero-copy architecture: shared data is not duplicated, consumers query provider's metadata and compute their own workloads.
- Works across regions and cloud platforms (AWS, Azure, GCP).
- Supports both private sharing (between accounts) and data marketplaces.

B. Security and Governance

- Strong Role-Based Access Control (RBAC) and policy-as-code support.
- Data masking, tokenization, and column-level security integrated with shares.
- End-to-end encryption and immutable audit logs.

C. Performance and Cost

- No storage duplication—cost savings for providers.
- Consumers pay only for compute resources.

- Near real-time sharing supports low-latency analytics across financial ecosystems.

D. Limitations

- Vendor lock-in to Snowflake ecosystem.
- Governance policies must be carefully aligned with regulatory frameworks.
- Initial setup complexity for multi-cloud replication.

IV. Comparative Analysis

Feature	Redshift Data Sharing	Snowflake Secure Data Sharing
Scope	Intra-AWS, limited cross-region	Cross-cloud, cross-region (Snowgrid)
Security	AWS IAM + KMS, Lake Formation	RBAC, masking, audit logs, policy-as-code
Compliance	Dependent on AWS compliance stack	Native GDPR/CCPA/FCRA-aligned features
Performance	Good for intra-account workloads	Optimized for real-time global sharing
Cost Model	Provider bears infra cost	Provider pays storage; consumer pays compute
Governance	AWS account-based	Fine-grained, tag-based, column-level

V. Use Cases in Financial Services

- Redshift Data Sharing: Best suited for financial institutions already standardized on AWS, requiring internal business unit collaboration.

-. Snowflake Secure Data Sharing is the **preferred solution** for institutions with **heterogeneous technology stacks**, **multi-cloud strategies**, and **global compliance obligations**. Its **zero-copy architecture** and **centralized governance** make it ideal for **financial ecosystems** that require **cross-entity collaboration** with strict control and auditability.

VI. Conclusion

Secure data sharing is central to financial services transformation. While Redshift Data Sharing provides a cost-effective, AWS-centric option, Snowflake Secure Data Sharing enables global, cross-cloud, compliance-ready collaboration. Institutions must evaluate based on regulatory requirements, ecosystem alignment, and collaboration scope. A hybrid approach—leveraging Redshift for intra-AWS workloads and Snowflake for external exchanges—can deliver the best of both worlds.



References

1. GDPR Compliance Overview, 2025. <https://gdpr.eu>
2. California Consumer Privacy Act (CCPA), 2025. <https://oag.ca.gov/privacy/ccpa>
3. AWS Redshift Documentation, 2024. <https://aws.amazon.com/redshift/>
4. Snowflake Documentation, 2024. <https://docs.snowflake.com>
5. U. Nayak, "Cost Optimization Strategies in Cloud Data Warehousing: AWS Redshift vs. Snowflake," *IJCEM*, vol. 8, no. 2, pp. 1–4, 2025.
6. U. Nayak, "Disaster Recovery in the Cloud: Best Practices for High Availability in Financial Services," *IJLRP*, vol. 6, no. 5, pp. 1–5, 2025.
7. U. Nayak, "Automated Data Governance and Compliance Monitoring using AI & Big Data," *IJRMPS*, vol. 13, no. 4, pp. 1–3, 2025.
8. U. Nayak, "Building a Multi-Cloud Data Strategy with AWS, Snowflake, and Terraform," *IJRMPS*, vol. 13, no. 5, pp. 1–4, 2025.
9. S. Gupta, "Fairness-Constrained Collections Policies: A Theoretical Model with Off-Policy Evaluation," *International Journal on Science and Technology*, vol. 16, no. 3, July–September 2025.