

SNOWFLAKE

SERVICE DEFINITION

DOCUMENT – LOT 2B

PREPARED FOR: G-CLOUD 15 FRAMEWORK

TABLE OF CONTENTS

SNOWFLAKE OVERVIEW	2
ARCHITECTURE	4
Snowflake Platform Architecture	4
Business Continuity	5
Snowflake Cost	6
DATA ARCHITECTURE PATTERNS WITH SNOWFLAKE	7
Data Lake/Lakehouse	8
Data Mesh	9
Applications	9
DATA INGESTION AND TRANSFORMATION	11
Openflow	11
Native Connectors	12
Snowpipe and Snowpipe Streaming	12
Partner ETL tools	12
Data Sharing	12
dbt Projects	14
Dynamic Tables	14
Streams and Tasks	15
Migrating to Snowflake	15
UNIFIED PLATFORM FOR END-TO-END AI, ML AND BI	16
Snowflake Cortex	16
Snowflake ML	19
Data Visualization and Reporting	21
SECURITY AND GOVERNANCE	23
Snowflake Horizon	23
Security and Governance Details	25
Snowflake AI Trust and Safety	26
CERTIFICATIONS AND COMPLIANCE	27
QUICKSTART PROGRAM FOR PUBLIC SECTOR	33

SNOWFLAKE OVERVIEW

This document sets out a high-level overview of the Snowflake service. For the avoidance of doubt, the full description of the Snowflake service is set out in the Snowflake Documentation available at <https://docs.snowflake.com/en/>.

The Snowflake AI Data Cloud Platform is a globally available platform for all data and AI workloads. Snowflake's cloud-native architecture was built from the ground up to be **easy to use, endlessly adaptable, interoperable, and future-proof**.

From its inception, Snowflake was designed to be a **fully managed**, unified software-as-a-service (SaaS) solution, eliminating the administrative overhead and infrastructure management of traditional platforms. This allows teams to focus on deriving value from their data rather than spending time on manual tuning and maintenance. The platform's guiding principles—from separating compute and storage to its multi-cloud, single-governance model—ensure a cost-effective and efficient experience as organizations grow.

Snowflake's **adaptable architecture** connects disparate data silos, whether they are on-premises or across different cloud providers like AWS, Azure, and GCP. It seamlessly handles **structured, semi-structured, and unstructured data**, offering **interoperable storage** that works with data in Snowflake's optimized format or in open table formats within the customers' own cloud storage, eliminating the need for complex and costly data movement processes.

As new technologies like **Generative AI** emerge, Snowflake integrates them directly into the platform, so that organizations can take advantage of the next disruption without re-architecting their environment. Snowflake designed the platform to be a connected ecosystem, facilitating **secure data sharing** and **collaboration** both internally and with external partners. By providing a single **secure, governed, and cost-effective** platform, Snowflake empowers organizations to build for today and adapt for tomorrow.

ARCHITECTURE

Snowflake Platform Architecture

Snowflake is the only cloud data platform with all components tightly integrated. Snowflake is delivered as a service, meaning that Snowflake takes care of managing the underlying cloud infrastructure, data protection, security and service tuning, so that you can focus on what's important: getting value from the data.

Responsibility Matrix: Snowflake vs Other Data Lakes & Data Warehouses				
	Task	Ordinary on-prem DW/DL	Ordinary cloud DW/DL	Snowflake
Infrastructure	Hardware	Customer	Cloud vendor	Cloud Data Platform as a Service
	Software	Distribution	Cloud vendor	
	Hardware cluster setup	Customer	Customer	
	Software provisioning	Tools	Tools	
Data & service protection	Data protection & retention	Platform or Add-on	Platform or Add-on	
	Node failure protection & recovery	Customer	Customer	
	Disaster recovery	Customer	Customer	
	Service monitoring & alerting	Customer	Customer	
Security	Physical security	Customer	Cloud vendor	
	Deployment security	Customer	Customer	
	Security monitoring	Customer	Customer	
Database management & tuning	Compute scaling	Customer	Customer	
	Index management	Customer	Customer	
	Data partitioning	Customer	Customer	
	Metadata & statistics management	Customer	Customer	
	Query optimization	Customer	Customer	

Snowflake implements a multi-cluster, shared data platform on Amazon Web Services, Microsoft Azure, and Google Cloud Platform. At its core, Snowflake is a columnar-store ANSI-SQL relational database, but its unique architecture enables a variety of workloads beyond traditional data warehousing, including lakehouse analytics, data engineering, AI, applications and collaboration.

Core components of the Snowflake architecture:

- Inside the Snowflake Platform architecture, the **Interoperable Storage** layer gives customers the choice whether to utilize Snowflake-managed or customer-managed cloud storage. When data is stored in Snowflake-managed storage, Snowflake automatically manages encryption, compression, partitioning, and clustering of the data in the storage layer. For customers who choose to store their data externally in open table formats, Apache Iceberg™ tables for Snowflake combine the performance and query semantics of typical Snowflake tables with external cloud storage that you manage.
- The **Multi-Cluster Compute** layer represents the independently scalable compute clusters that process queries and updates against the underlying Snowflake database. Snowflake compute clusters elastically scale vertically, horizontally and across workloads and are configured by default to shut down when demand falls to zero for a specified period of time. The **Adaptive Compute** service automatically adjusts compute resources based on the workload demands. Snowflake's compute can execute workloads in SQL, Python, Java, Scala and Javascript, so there is no need to move the work and data between different compute engines. Both CPU and GPU compute resources are available based on the workload demands.
- **Cortex AI** delivers a fully integrated layer of Generative AI services, including state-of-the art LLM models hosted on scalable GPU infrastructure, and easy-to-use tools to leverage them inside your pipelines and workflows. Use **AI SQL** for powerful multimodal LLM-based analytics directly inside your SQL queries; process documents at scale with **Document AI**; leverage **Snowflake Intelligence** for AI agent-powered conversational analytics over your structured and unstructured data with unparalleled accuracy.
- The **Cloud Services** layer sits over the storage, compute, and AI layers to manage and optimize the platform's operations. As workloads are added or removed from the Snowflake platform, the cloud services layer automatically scales to support optimization, security, transactional consistency, and metadata management with near-zero administration required. All you have to worry about is designing, loading, and querying your data, and Snowflake takes care of the rest.
- Finally, the **Snowgrid** is a cross-cloud technology layer that allows you to distribute your data and operations across the regions and cloud providers of your choice while maintaining the same Snowflake experience regardless of complexity. The Snowgrid enables a **multi-cloud strategy**, including a cross-cloud approach to failover and disaster recovery for unmatched business continuity, and to mix and match clouds as you see fit. Powered by Snowgrid, **Snowflake Marketplace** provides instant access to thousands of datasets and applications that organizations rely on – the most relevant and valuable content is ready to be mobilized to drive better decisions, products and services, and business outcomes.

Business Continuity

The Snowflake Data Cloud was designed from the ground up to be a continuously available cloud service, resilient to failures, disruption, or data loss. Its continuous data protection (CDP) capabilities protect against (and provide easy self-service recovery from) accidental errors, system failures, and malicious acts.

- **Customer error.** To mitigate accidental data deletion or corruption, Snowflake **Time Travel** enables accessing historical data (i.e. data that has been changed or deleted) at any point

within up to 90 days. Historical data can be retrieved and restored with a simple SQL query.

Fail-safe provides an additional 7 days period during which historical data may be recoverable by Snowflake.

- **Infrastructure failures** – from a single instance of virtual machine failing to an Availability Zone outage. The Snowflake architecture includes **triple redundancy** of all critical services, including built-in redundancy across **three availability zones** within a cloud region.
- **Regional or multi-region failures.** Customers with critical workloads can set up **cross-region/cross-cloud replication** and **failover/failback** including **client redirect** with just a few SQL commands - no need to set up complex processes and HA/DR infrastructure.

Snowflake Cost

Snowflake is a fully managed and fully integrated service. With other solutions, customers receive separate bills for the data warehouse/data lake service and for the cost of the underlying infrastructure such as storage, compute, file access and support. Snowflake customers receive **one bill**, with all costs billed directly by Snowflake.

The total cost of using Snowflake is the aggregate of the cost of using compute resources, storage and optional services such as snowpipe and materialized views. There are **no data egress charges** with Snowflake as long as customers use one of the many drivers and connectors to access the data and do not export files directly from Snowflake to another cloud or region object storage. Snowflake Support is also included in the cost of a Snowflake credit. To learn more, visit <https://docs.snowflake.com/en/user-guide/cost-understanding-overall>.

Snowflake provides many built-in cost management capabilities to help customers understand, control and optimize costs:

1. **Visibility** – understand, attribute, monitor spend:
 - The **Cost Management UI** provides an account-level consumption and spend, consumption trend for virtual warehouses, most expensive queries and more.
 - Customers can set up account- or project-specific **budgets** that define a monthly spending limit for a group of Snowflake objects. When the system projects the spending limit will be exceeded, Snowflake sends a daily notification email to the Snowflake admin.
 - Snowflake automatically monitors consumption trends for anomalies and can be set up to send alerts when unusual activity is detected.
 - Customers can query the usage views to drill down into cost data and generate custom reports and dashboards. The **Query Insights** view provides actionable recommendations for query optimization - such as pointing out an incomplete join condition generating a cartesian join, or spillage due to an undersized compute resource.
2. **Control** – Limit and control spend.
 - Snowflake virtual warehouse (compute) sizes provide implicit upper bounds for the credit burn rate and cost. In addition, virtual warehouses **auto-suspend and auto-resume** by default so customers only pay for what they use.
 - **Resource Monitors** provide alerting and hard limits for individual virtual warehouses during a specific time interval. The virtual warehouse can be configured to suspend once the credit limit is hit.

3. **Optimization** – Identify inefficiencies and optimize spend.
- Since Snowflake's founding in 2012, we strive to be the market leader in fast performance for all workloads, passing the optimizations to the customers automatically and increasing the value of a Snowflake credit. The **Snowflake Performance Index** measures the improvement in query runtime and other metrics based on these improvements. To learn more, visit <https://www.snowflake.com/en/pricing-options/performance-index/>.
 - Examples of optimization features include **Snowpark-optimized warehouses**, which reduce the time and cost of memory-intensive workloads; serverless capabilities such as serverless tasks, query acceleration service, automatic clustering, materialized views and search optimization, which deliver performance improvements for specific query patterns and help save on costs over time.
 - Recently introduced **Generation 2 (Gen2) virtual warehouses** run on upgraded hardware and deliver 2.1x faster performance YoY, with 2.3x higher throughput for Concurrent BI workloads and up to 4.4x faster Delete, Update and Merge operations.

DATA ARCHITECTURE PATTERNS WITH SNOWFLAKE

Organizations in every industry acknowledge that data as one of their most important assets. However, many are falling short of realizing the potential of their data for three common reasons:

- Even with the increasing adoption of the cloud to store data at virtually infinite volumes, organizations still struggle to centralize their data. Silos are everywhere—architectures, clouds, business entities, languages, data types, and table formats.
- To account for these silos, data engineers are tasked with creating complex, error-prone pipelines. An organization may be able to store massive amounts of data in their data lake, but processing and querying that data is still either slow, expensive, or both. And that's just first-party data. Incorporating shared, third-party data can be just as difficult and expensive: requiring making a copy to ETL to your data lake.
- And with lots of systems and lots of copies, governance is a nightmare and collaboration is difficult. This presents risk for compliance, and leads to unknown or poor data quality. Users can't easily find data, trust it's what they need, and then share to collaborate.

Snowflake's **adaptable architecture** allows you to truly manage all data together and unify the silos:

- Cross-region, cross-cloud, cross-entity
- Across programming languages
- Across structured & unstructured data
- Across compute engines and analytical/transactional workloads


Having flexibility in architectural patterns is critical to not only address organizations' business needs today but also will let them adapt to changing needs in the future. With Snowflake, you get one platform that can support a wide range of architectural patterns, including:

- Data warehouse: Centralized analytics for structured business data
- Data lake: Unlimited storage for versatile data types and workloads
- Data lakehouse: Transactional data lake for unified analytics, AI/ML, collaborative workloads
- Data mesh or fabric: Distributed and governed, domain-oriented collaboration
- Transactional applications: Build and run applications directly on Snowflake with Unistore's

Hybrid Tables or leverage with Snowflake PostgreSQL for low-latency, transactional needs.

Data Lake/Lakehouse

The purpose of data lakes is to liberate business data from applications that generate it, and make it widely available for analytics. The standard pattern replicates data from transactional systems into a raw or staging area in Snowflake and uses Snowflake's compute power to transform and enrich the data to create curated tables and data marts that end users use to run reports and create predictive models. This pattern delivers the most benefits for best performance, security, and automatic management. Snowflake-managed storage offers a reduced total cost of ownership, better time to value, and limits the risk of undesired data exposure and compliance risks. With Snowflake managing the secure data layer, customers can focus on business objectives over a myriad of technical decisions (e.g. partition size, encryption, access control).



Snowflake Horizon Catalog is embedded in the Snowflake service and provides out-of-the-box transactional consistency, unified security, governance and discovery across regions and clouds. Whenever you load data or run queries in Snowflake, Horizon operates behind the scenes without any configuration or intervention from you.

Some organizations prioritize **interoperability** with other compute engines such as Apache Spark, Flink and Trino and therefore don't mind the additional administrative burden of managing their data in customer-managed storage (e.g. partition size, encryption, access control, etc). For such customers,

Snowflake's extensive support of Apache Iceberg™ table format delivers this interoperability. Iceberg tables for Snowflake combine the performance and query semantics of typical Snowflake tables with external cloud storage that you manage.

Horizon Catalog federates across Iceberg REST catalogs (including the open source **Apache Polaris** catalog) through **Catalog-linked Databases**, which automatically syncs Iceberg objects managed by remote catalogs in Horizon Catalog. Data engineers can act on their entire Iceberg ecosystem from a single pane of glass, ensuring all read and write operations stay in sync and are protected by Horizon Catalog's leading governance and discovery capabilities.

Data Mesh

Snowflake's platform provides a robust foundation for implementing a **data mesh** architecture by directly supporting its core principles. It enables **domain-oriented ownership** by allowing individual business domains to manage their data and infrastructure independently within a single, unified platform. The cornerstone of its support is **Secure Data Sharing**, which facilitates the creation of **data as a product**. This feature allows domains to share live, ready-to-query data directly from their accounts with other domains without creating complex ETL pipelines or data copies, ensuring consumers always have access to fresh, governed data.

Furthermore, Snowflake enables **federated computational governance** through features like dynamic data masking, row-access policies, and object tagging. These tools empower domain teams to apply fine-grained security and governance controls locally on their data products while adhering to global organizational standards. With the integration of **Snowpark** and **Streamlit in Snowflake**, data products can extend beyond simple datasets to include interactive applications and complex data processing logic, all discoverable and accessible within the same ecosystem. This combination of secure, live sharing and decentralized governance makes Snowflake an effective platform for building and scaling a data mesh.

Applications

For decades, transactional and analytical data have remained separate, significantly limiting how fast organizations could evolve their businesses. Snowflake is changing this by unifying these disparate worlds, allowing you to build and run applications directly on the same platform as your analytics. This transformation is powered by a suite of foundational innovations: **Unistore** for hybrid transactional-analytical tables, a fully managed **Postgres service** (based on the Crunchy Data acquisition) for OLTP applications, **Snowpark Container Services** for running any containerized logic, and the **Snowflake Native App Framework** to package and deliver these new solutions securely to customers.

With **Unistore**, organizations can use a single, unified data set to develop and deploy applications, and analyze both transactional and analytical data together in near-real time. Underpinning the Unistore workload is a **hybrid table** - a Snowflake table type that is optimized for hybrid transactional and operational workloads that require low latency and high throughput on small random point reads and writes. Hybrid tables accelerate application development by accessing analytical and transactional data with a single query and provide the following key benefits:

- **Fewer Databases:** Hybrid tables run double-digit millisecond point operations alongside your analytical queries, all within Snowflake. Reduce the need to operate and synchronize across

multiple databases. With a single governed dataset for both your transactional and historical data, your application data is ready to query, no ETL required.

- **Get started quickly:** Query hybrid tables with the syntax you are used to in Snowflake. Hybrid tables enforce primary keys and referential integrity constraints, as well as provide secondary indexes for accelerated lookups. Snowflake governance capabilities like masking policies and row based access policies work out of the box.
- **Easily bring transactional data to your analytical workloads:** With your transactional and analytical data in a single platform, you can enrich your analytics and create new experiences.

In addition to Unistore, Snowflake is in the process of building a fully **managed Postgres** service based on the recent acquisition of **Crunchy Data**, a leading open source Postgres provider. This service is engineered to run OLTP workloads natively within the Snowflake AI Data Cloud and will feature full wire-protocol compatibility with open-source PostgreSQL. This compatibility is critical, as it allows developers to migrate existing applications with minimal to no code refactoring. Architecturally, the service facilitates seamless, high-performance data interoperability between the managed Postgres instances and Snowflake's core analytical engine, eliminating the need for separate ETL pipelines. Crucially, this integration extends Snowflake's **unified governance and security** model to Postgres, allowing a single control plane for applying granular access policies, data masking, and object tagging across both transactional and analytical data estates.

The **Snowflake Native App Framework** is a developer framework that enables providers to build, distribute, and monetize Snowflake Native Apps on Snowflake Marketplace. Snowflake Native Apps run inside the end-consumer's Snowflake account.

Using the framework, providers can build applications with Snowflake functionalities like UDFs, stored procedures, and Snowpark. Functionalities have been extended with the capability to provide low-code UX framework through the Streamlit Integration and to empower providers to craft more complex apps via the integration with Snowpark Container Services.

Providers can deploy and operate Snowflake Native Apps with out-of-the-box versioning and monitoring support, distribute to Snowflake Marketplace through public or private listing, and leverage billing infrastructure to monetize apps using customizable billing models.

Consumers can purchase apps with a credit card or their Snowflake credit, install and run these apps within their Snowflake accounts, running directly against their data securely in the Data Cloud.

Snowpark Container Services is a fully managed container offering designed to facilitate the deployment, management, and scaling of containerized applications within the Snowflake ecosystem. This service enables users to run containerized workloads directly within Snowflake, ensuring that data doesn't need to be moved out of the Snowflake environment for processing. Unlike traditional container orchestration platforms like Docker or Kubernetes, Snowpark Container Services offers an OCI runtime execution environment specifically optimized for Snowflake. This integration allows for the seamless execution of OCI images, leveraging Snowflake's robust data platform.

As a fully managed service, Snowpark Container Services streamlines operational tasks. It handles the intricacies of container management, including security and configuration, in line with best practices. This ensures that users can focus on developing and deploying their applications without the overhead of managing the underlying infrastructure. Snowpark Container Services is fully integrated with Snowflake and with third-party tools like Docker. Seamless integration makes it easier for teams to focus on building data applications.

Snowpark Container Services unlocks a wide array of new functionality, including these features:

- Create long-running services.
- Use GPUs to boost the speed and processing capabilities of a system.
- Write your application code in any language (for example, C++).
- Use any libraries with your application.

The integration of Snowpark Container Services with the Snowflake Native App Framework fundamentally expands the types of applications providers can build and monetize. It allows developers to package complex logic, written in any language (e.g., C++, Python, Java) and using any library, into a container. This container can then be deployed as part of a Native App, enabling features like long-running services or GPU-accelerated processing to run directly and securely against a consumer's data within their own Snowflake account. All of this comes with Snowflake platform benefits, most notably ease-of-use, security, and governance features. And you now have a scalable, flexible compute layer next to the powerful Snowflake data layer without needing to move data off the platform.

DATA INGESTION AND TRANSFORMATION

Data ingestion & transformation is a key component of a robust data platform. A data platform has limited value if you can't bring in the data. Snowflake provides an extensive array of data ingestion and transformation services, giving you the choice and convenience to build reliable, scalable data pipelines tailored to any source, format, or latency requirement.

Openflow

Snowflake Openflow is a new **managed data integration service** within the Snowflake platform, built on Apache NiFi, designed to connect and process data from **virtually any source** into Snowflake and other destinations. Openflow unifies the entire data pipeline — from ingestion of batch, streaming, and unstructured data to transformation — within a single, governed platform, eliminating the need for separate ETL tools and accelerating the development of AI applications.

Key benefits of Openflow:

- **Open and extensible:** An extensible managed service that's powered by Apache NiFi, enabling you to build and extend processors from any data source to any destination.
- **Unified data integration platform:** Openflow enables data engineers to handle complex, bi-directional data extraction and loading through a fully managed service that can be deployed inside your own VPC in the cloud or on-premises.
- **Enterprise ready:** Openflow offers out-of-the box security, compliance, and observability and maintainability hooks for data integration.

- **High speed ingestion of all types of data:** One unified platform that lets you handle structured and unstructured data, in both batch and streaming modes from your data source to Snowflake at virtually any scale.
- **Continuous ingestion of multimodal data for AI processing:** Near real-time unstructured data ingestion, so you can immediately chat with your data coming from sources such as Sharepoint, Google Drive, and so on.

Architecturally, Openflow is comprised of:

- **Control plane:** runs in Snowflake as a managed service, allowing customers to provision, command, control data plane.
- **Data Plane:** Openflow connectors and pipelines run in the data plane environment. You can choose to run the Dataplane in your own VPC via Bring Your Own Cloud (BYOC) or in Snowflake as a fully managed service. Data planes can scale horizontally, support multi-node and multi-cluster deployments, and DR resilience. You can have multiple data planes tied to the control plane, serving different use cases.
- **Openflow Connectors:** Highly curated, fully managed, first-class, secure, governed experiences against key systems (unstructured data, SaaS systems, databases, streaming etc). Connectors are extensible - you can create custom connectors and run them in the data plane.
- **OSS Nifi:** Openflow is powered by the open source Nifi data integration framework that is trusted by thousands of enterprise customers. Openflow super charges Nifi.

Native Connectors

Snowflake and partner-delivered connectors to a variety of data sources can be provisioned from the Snowflake Marketplace so that you can start replicating data in minutes.

Snowpipe and Snowpipe Streaming

Snowflake supports low-latency, continuous data ingestion via micro-batch ingestion (using Snowpipe) and streaming (using Snowpipe Streaming). Data is ingested automatically from sources such as cloud object storage (AWS S3, Azure blob, Google Cloud Storage) and streaming platforms like Kafka. Loading happens continuously as new data is generated, eliminating the need for scripts and scheduling tools. The COPY INTO command can be used to perform the initial bulk load.

Partner ETL tools


Snowflake supports a vast community of extract, transform, and load tools (ETLs) to streamline the continuous movement of data into and out of the Snowflake cloud-native platform. Companies such as Fivetran, Matillion, and Informatica have created tight integrations with Snowflake through dedicated connectors. See the Snowflake ecosystem for more details: <https://docs.snowflake.com/en/user-guide/ecosystem>.

Data Sharing

By far the easiest way to get data into Snowflake is through secure data sharing between accounts of Snowflake. With secure data sharing, no actual data is copied or transferred between accounts. This unique architecture provides the foundation for privacy-preserving collaboration, as organizations can grant secure access to live data while maintaining full control and applying governance. Because no data is copied or exchanged, secure data sharing setup is quick and easy for providers and access to


the shared data is near-instantaneous for consumers. You can share data, semantic views, AI models, and applications.

Data Clean Rooms leverage Snowflake's data sharing capabilities and allow you to combine and analyze data from different parties with privacy-preserving configurations that help protect the underlying data. Delivered as a native solution, clean rooms makes it easy to build, connect, and use data from multiple parties, while applying enhanced privacy and increased security.



Snowflake Marketplace - another feature that leverages Snowflake's data sharing capabilities - provides access to data products from hundreds of data providers for categories ranging from weather to demographics to economics to healthcare and energy.

The **Internal Marketplace** in Snowflake is a curated, secure space for collaborative data sharing within an organization. It centralizes internally-available data products, allowing teams to easily discover and use vetted data shared within your internal teams. Access can be managed by account targeting and Role-Based Access Control (RBAC), ensuring that data remains secure and accessible only to authorized users.



dbt Projects

dbt Core is an open-source data transformation tool and framework that many customers use to define, test, and deploy transformations. A favorite among data teams, dbt offers the ability to bring software engineering best practices and improved efficiency in SQL and Snowpark data transformation workflows right in Snowflake. **dbt projects in Snowflake** enable data teams to build, run and monitor dbt projects directly in Snowsight UI, reducing context switching, simplifying setup and accelerating the development lifecycle of data pipelines.

With dbt Projects support, you can:

- Enable new teams to build and deploy pipelines with uniform governance: Accelerate onboarding and empower new teams to create the pipelines they need through an intuitive interface with a uniform governance and security model for data and pipelines.
- Consolidate systems to reduce administration and improve debugging: Run dbt natively on Snowflake and streamline the pipeline development lifecycle to improve developer productivity and quickly spot and address problems.

Dynamic Tables

As the raw data lands in Snowflake, **dynamic tables** are the building blocks of declarative data transformation pipelines. They significantly simplify data engineering in Snowflake and provide a reliable, cost-effective, and automated way to transform your data incrementally using embedded change data capture (CDC). Instead of defining data transformation steps as a series of tasks and then

monitoring dependencies and scheduling, you can simply define the end state of the transformation using dynamic tables and leave the complex pipeline management to Snowflake.

Streams and Tasks

Snowflake **streams** and **tasks**, in combination with **stored procedures** and **UDFs**, enable you to design sophisticated data pipelines, particularly when the logic requires the use of Python, Java or Scala libraries.

The Snowflake UI makes monitoring data transformation pipelines easy by providing graphical observability interfaces displaying the pipelines, their status, dependencies and refresh history.



Migrating to Snowflake

Traditional data platforms and big data solutions often fail to achieve their main goal: allowing users to work with data without restrictions on size, speed, or adaptability. Snowflake provides a robust migration framework designed to help organizations seamlessly transition from legacy data platforms and modernizing their data analytics workloads.

SnowConvert is a user-friendly tool that helps you modernize your traditional data platform by migrating it to Snowflake. SnowConvert analyzes your existing SQL code (from platforms like Oracle, SQL Server, and Teradata) and converts it to Snowflake SQL. After conversion, you can immediately take advantage of all Snowflake's features and capabilities.

Currently, SnowConvert supports converting code from these source platforms:

- Teradata
- Oracle
- SQL Server
- Sybase IQ
- Redshift
- Azure Synapse

While SnowConvert automates a high percentage of the initial code conversion, some complex or non-

standard code may still require manual review. **SnowConvert AI** extends the power of the core tool by tackling this final stage of the migration.

Integrated into Visual Studio Code as the "SnowConvert Migration Assistant," it acts as an AI-powered assistant that helps developers resolve the issues that SnowConvert could not automatically fix. Powered by **Snowflake Cortex AI**, it intelligently analyzes the remaining code, providing natural language explanations for conversion problems and suggesting specific code fixes. By streamlining this often complex and manual cleanup process, SnowConvert AI dramatically reduces the overall complexity, costs, and timelines associated with data warehouse, BI, and ETL migrations, making the journey to the Snowflake AI Data Cloud more accessible than ever.


UNIFIED PLATFORM FOR END-TO-END AI, ML AND BI

Snowflake delivers a comprehensive, end-to-end platform that supercharges AI and ML by unifying its powerful capabilities into a single, fully managed service. The platform combines **Snowflake ML**, which supports the complete predictive modeling lifecycle (from feature engineering and model training to deployment and monitoring) with **Snowflake Cortex AI**, which provides ready-to-use, fully managed LLM models and features that make it easy to deploy generative AI applications on your data. Because these capabilities are fully managed and deeply integrated, customers can build, deploy, and govern everything from custom models to advanced AI agents directly on their data without the complexity of moving data or stitching together separate tools. This unified approach is built on three core value propositions: it's **easy**, reducing development costs by eliminating infrastructure management; it's **connected**, enabling better quality and operational efficiency by having all components work together seamlessly; and it's **trusted**, extending Snowflake's industry-leading security and unified governance to all data and models within one secure perimeter.

In addition, Snowflake delivers native data exploration and visualization experiences via **Workspaces** and **Streamlit**. For enterprise BI, Snowflake seamlessly integrates with third-party BI solutions such as Power BI, Tableau, Sigma, Quicksight and more.

Snowflake Cortex

The world of business analytics is experiencing its most significant transformation in generations. Generative AI is at the heart of this evolution, fundamentally democratizing enterprise AI and redefining user experiences. The focus is shifting from complex dashboards to intelligent, conversational AI agents that allow anyone to ask questions in natural language and get trusted answers. Snowflake is leading this charge with an **easy, efficient, and trusted** platform where these capabilities are fully managed and integrated. **Snowflake Cortex** - a fully managed AI layer within Snowflake's secure perimeter - enables business users and data scientists to turn structured and unstructured data into actionable insights, without complex tooling or infrastructure. Whether customers are analyzing tables, processing documents, deploying AI agents, or training models, Snowflake integrates these into its secure platform with unified governance, eliminating the need for infrastructure management or complex tooling.



Snowflake Intelligence provides a unified conversational AI interface integrated directly within the Snowflake platform, enabling business users to analyze, summarize, and act on both structured and unstructured data using natural language. This interface eliminates the need for traditional dashboards or SQL, allowing business users to access and derive insights easily and reliably. For developers, it reduces technical friction by providing the flexibility to quickly customize experiences, integrate new data sources, and expand capabilities, ultimately helping them serve business users more efficiently and effectively.

Cortex Agents power the conversational experiences within Snowflake Intelligence by coordinating tasks across both structured and unstructured data sources to provide valuable insights. Cortex Agents plan their tasks, utilize various tools to execute them, and generate responses. They can leverage Cortex Analyst, Cortex Search, and large language models (LLMs) as tools for data analysis.

Cortex Agents use Cortex Analyst and Cortex Search as tools, along with LLMs, to analyze data. Cortex Search extracts insights from unstructured sources, while Cortex Analyst generates SQL to process structured data. A comprehensive support for tool identification and tool execution enables delivery of sophisticated applications grounded in enterprise data.

REST API make it easy to call Cortex Agents from external applications, and integration with Microsoft Teams and M365 Copilot allows users to fit the agents directly into their daily workflow, while staying within Microsoft and Snowflake's trusted security boundaries and working with your existing semantic models and data governance policies.

Cortex Analyst is a fully-managed, LLM-powered feature that helps you create applications capable of reliably answering business questions based on your **structured data** in Snowflake. With Cortex Analyst, business users can ask questions in natural language and receive direct answers without writing SQL.

Building a production-grade conversational self-service analytics solution requires a service that generates accurate text-to-SQL responses. For most teams, developing such a service that


successfully balances accuracy, latency, and costs is a daunting task. Cortex Analyst simplifies this process by providing a fully managed, sophisticated agentic AI system that handles all of these complexities, generating highly accurate text-to-SQL responses. It helps you accelerate the delivery of high-precision, self-serve conversational analytics to business teams, while avoiding time sinks such as complex RAG solution patterns, model experimentation, and GPU capacity planning. The generated SQL queries are executed against the scalable Snowflake engine, ensuring industry-leading price performance and lower total cost of ownership (TCO).

Generic AI solutions often struggle with text-to-SQL conversions when given only a database schema, as schemas lack critical knowledge like business process definitions and metrics handling. Cortex Analyst overcomes this limitation by leveraging **semantic models**, which define business metrics and model business entities and their relationships. Cortex Analyst reads the information captured in the semantic model definition and generates the SQL against the physical tables directly.

Automatic semantic model generation makes creating semantic models for Cortex Analyst easier and faster. By analyzing schema metadata, query history and dashboards, it constructs performant, reusable models, eliminating weeks of manual work.

Cortex Search is a feature that uses state-of-the-art models for fast retrieval and ranking of information from **unstructured** documents. Cortex Search enables low-latency, high-quality “fuzzy” search over your Snowflake data. Cortex Search powers a broad array of search experiences for Snowflake users including Retrieval Augmented Generation (RAG) applications leveraging LLMs.

Cortex Search gets you up and running with a hybrid (vector and keyword) search engine on your text data in minutes, without having to worry about embedding, infrastructure maintenance, search quality parameter tuning, or ongoing index refreshes. This means you can spend less time on infrastructure and search quality tuning, and more time developing high-quality chat and search experiences using your data. Check out the Cortex Search tutorials for step-by-step instructions on using Cortex Search to power AI chat and search applications.




AISQL

Cortex AISQL embeds AI capabilities directly into SQL environments, introducing native AI operators like `AI_FILTER`, `AI_JOIN`, and `AI_AGG` within the Snowflake ecosystem. This allows users to build sophisticated analysis pipelines across multiple data types (text, image, and soon audio) without needing a separate AI stack. It empowers data analysts to perform advanced tasks, such as semantic joins and multi-label classification, all within their familiar SQL interface.

Document AI is a fully managed service that uses the powerful **Arctic-TILT**, a proprietary large language model to automatically extract information like text, table values, and handwritten content from documents such as PDFs. It is designed to improve efficiency and reduce the costs and errors associated with manual data entry. Business users can take advantage of a no-code interface to test and refine the AI's extraction capabilities, while data teams can use SQL to extract data from multiple documents at once and load it directly into Snowflake tables.

Documents review



Manual_2022-02-01.pdf

Values to extract

machine

what is the machine that was inspected?

1:00 Injection Molder

serial_number

what is the serial number of the machine?

1:00 SGMM-12345

inspection_grade

what is the overall inspection grade?

1:00 PASS

failed_machine_parts

are there any machine parts that failed?

1:00 No

machine_parts

what are the machine parts that were inspected?

1:00 Injection Unit

1:00 Mold Clamping Unit

1:00 Hydraulic System

1:00 Temperature Control System

1:00 Ejector System

1:00 Lubrication System

1:00 Safety Devices

1:00 Control Software

1 / 1

100%

Accept all and review next


Snowflake ML

Snowflake ML is an integrated set of capabilities for end-to-end machine learning in a single platform on top of your governed data. This is a unified environment for ML development and productionization that is optimized for large-scale distributed feature engineering, model training and inference on CPU and GPU compute without manual tuning or configuration.

Scaling end-to-end ML workflows in Snowflake is seamless. You can do the following:

- Prepare data
- Create and use features with the Snowflake Feature Store
- Train models with CPUs or GPUs using any open-source package from Snowflake Notebooks on Container Runtime.
- Operationalize your pipelines using Snowflake ML Jobs.
- Deploy your model for inference at scale with the Snowflake Model Registry
- Monitor your production models with ML Observability and Explainability.
- Use ML Lineage provides to track the source data to features, datasets, and models throughout the ML pipeline.

End-to-End ML Workflow with Snowflake ML



Snowflake ML is also flexible and modular. You can deploy the models that you've developed in Snowflake outside of Snowflake and externally-trained models can easily be brought into Snowflake for inference.

Snowflake Notebooks on Container Runtime provide a Jupyter-like environment for training and fine-tuning large-scale models in Snowflake, without infrastructure management. Start training with preinstalled packages such as PyTorch, XGBoost, or Scikit-learn, or install any package from open-source repositories like HuggingFace or PyPI. Container Runtime is optimized to run on Snowflake's infrastructure to provide you with highly efficient data loading, distributed model training, and hyperparameter tuning.

Snowflake's **Data Science Agent** represents a significant advancement in machine learning workflow automation, designed to enhance the productivity of data scientists through natural language interactions. This innovative agentic companion streamlines the entire machine learning pipeline by automating crucial steps that traditionally require extensive manual coding and configuration. The agent

assists data scientists by facilitating data preparation processes, automating feature engineering tasks, and streamlining model training workflows through intuitive natural language commands.

The Snowflake Feature Store is an integrated solution for defining, managing, storing and discovering ML features derived from your data. The Snowflake Feature Store supports automated, incremental refresh from batch and streaming data sources, so that feature pipelines need be defined only once to be continuously updated with new data.

Use **Snowflake ML Jobs** to develop and automate ML pipelines. ML Jobs also enable teams that prefer working from an external IDE (VS Code, PyCharm, SageMaker Notebooks) to dispatch functions, files or modules down to Snowflake's Container Runtime.

The **Snowflake Model Registry** allows for the logging and management of all your ML models, regardless of whether they're trained on Snowflake or other platforms. You can use the models from the model registry to run inference at scale. You can use Model Serving to deploy the models to Snowpark Container Service for inference.

ML Observability provides tools to monitor model performance metrics in Snowflake. You can track models in production, monitor performance and drift metrics, and set alerts for performance thresholds. Additionally, use the ML Explainability function to compute Shapley values for models in the Snowflake Model Registry, regardless of where they were trained.

ML Lineage is a capability to trace end-to-end lineage of ML artifacts from source data to features, datasets, and models. This enables reproducibility, compliance, and debugging across the full lifecycle of ML assets.

For business analysts, use **ML Functions** to shorten development time for common scenarios such as forecasting and anomaly detection across your organization with SQL.


Snowflake AI & ML Studio provides intuitive graphical interfaces to many of the Snowflake Cortex and Snowflake ML components.

Data Visualization and Reporting

As conversational AI forges new paradigms for dynamic data interaction, it complements, rather than replaces, the critical business requirement for structured reporting to ensure consistent, governed analysis and performance monitoring.

Snowflake tackles this requirement with a two-pronged approach. For custom needs, its native **Streamlit** integration allows developers to rapidly build interactive applications and visualizations that run securely within Snowflake, directly on live data. For established BI workflows, its architecture supports high-performance connectors to all leading third-party tools, ensuring that even complex dashboards remain fast and responsive.

Streamlit is an open-source Python library that makes it easy to create and share custom web apps for machine learning and data science. By using Streamlit you can quickly build and deploy powerful, customizable data applications. **Streamlit in Snowflake** helps developers securely build, deploy, and share Streamlit apps on Snowflake's data cloud. Using Streamlit in Snowflake, users can build applications that process and use data in Snowflake without moving data or application code to an external system, or worrying about the infrastructure that the applications run on.



Industry-standard BI tools connect to Snowflake via **native connectors** and standard **ODBC/JDBC drivers**, enabling users to leverage meaningful insights at scale with the tools they already have or with the latest generation of visualization technologies.

Sigma, Thoughtspot, Tableau and Power BI and many other popular BI tools integrate tightly with Snowflake and support high performance reporting on live data. Thousands of Snowflake customers leverage the power and scalability of Snowflake using these tools without the need to write a single line of code.

SECURITY AND GOVERNANCE

Snowflake Horizon

Snowflake Horizon is Snowflake's built-in governance solution with a unified set of compliance, security, privacy, interoperability, and access capabilities. Snowflake Horizon makes it easy for customers to govern and take immediate action on data, apps, and more across clouds, teams, partners, and customers — both inside and outside of organizations. Copilot for Horizon Catalog (private preview soon) will allow you to use natural language to perform governance, security and metadata discovery tasks in Snowsight, simplifying the process of getting critical information about data sets. Integrations with other Apache Iceberg-compatible catalogs and engines, and with partners, further extend Snowflake Horizon across customers' ecosystems.

Snowflake Horizon protects and audits customers' data with compliance, business continuity, data quality monitoring, and lineage. Snowflake maintains an extensive number of [compliance certifications](#) and continues to add more.

Protect and audit your data with compliance, business continuity, data quality monitoring and lineage:

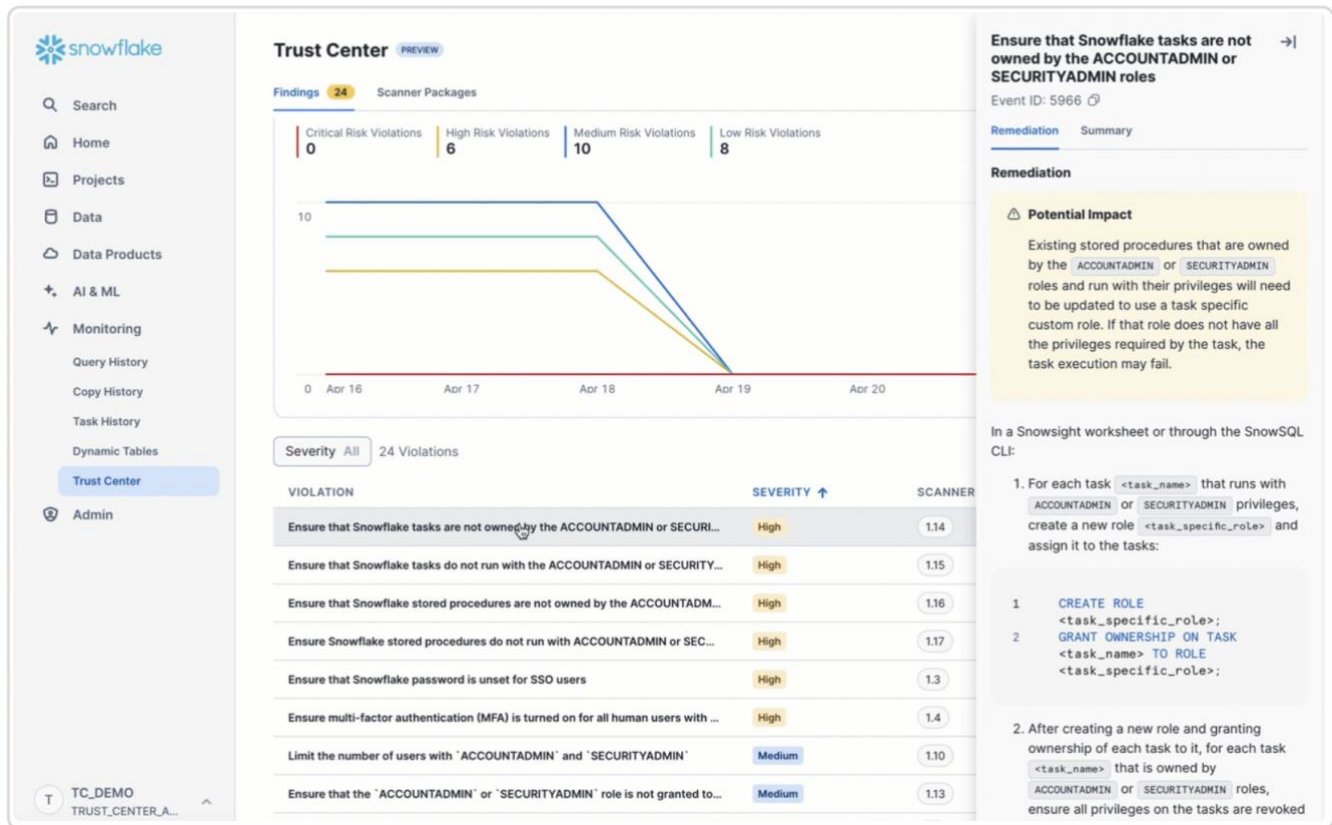
- Secure data on a platform with comprehensive compliance certifications
- Replicate and failover databases, shares, role-based access control, governance policies, compute resources, network policies, and more across clouds
- Monitor data quality with both out-of-the-box and custom metrics
- Audit data usage through Access History and Schema Change Tracking
- Understand the flow of data through object dependencies and lineage

Security

Secure your environment with continuous risk monitoring and protections, RBAC and granular authorization policies:

- [Trust Center](#) allows customers to monitor and manage risks 24/7 across your Snowflake accounts and cross cloud service providers. The Trust Center evaluates your account against recommendations specified in [scanners](#) according to a schedule you control. If your account violates any of the recommendations in any of the enabled scanners, then the Trust Center provides a list of security risks, and information about how to mitigate those risks.
- Safeguard data with built-in platform security capabilities across areas like authentication, encryption, connectivity and networking
- Enforce row and column-level security across clouds with RBAC and policies, eliminating the need to manage multiple versions of the same data

Horizon Security: Trust Center




Privacy

Unlock the value of your sensitive data with advanced privacy policies and data clean rooms:

- Identify and tag PII and other sensitive data.
Automatically detect and tag sensitive data via [automatic sensitive data classification](#), which continuously monitors tables within a specific schema and classifies their columns using native and custom classification categories.
- Easily collaborate on sensitive data with multiple teams inside and outside your organization with Snowflake Data Clean Rooms.

Horizon: Sensitive Data Classification



fname	gen	age	phone
Jane Doe	F	50	333-555-1236
Michael Gaines	M	75	666-666-1357
Ann Marshall	F	48	555-555-1234
John Smith	M	39	123-555-1234

Interoperability

Integrate with other Apache Iceberg-compatible catalogs and engines, in addition to data catalog and data governance partners:

- Iceberg Tables can be configured to use either Snowflake or an external service as the catalog to track metadata.
- Pre-built Snowflake integrations by leading data governance, security and data observability data catalog and data governance partners allow customers to maximize existing data investments and better manage their entire data estate.
- Horizon Catalog will provide external data discovery (private preview soon) for external data in relational databases, dashboards, semantic models and more.

Access

Classify, share, discover and immediately act on data, apps and more across regions and clouds:

- Identify, tag, and track sensitive data.
- Discover data and apps within and across organizations.
Use AI to [automatically generate descriptions](#) for a column, table, or view. The Cortex-Powered Object Descriptions feature leverages Snowflake-hosted large language models (LLMs) to evaluate object metadata and, if desired, sample data to generate the description.
- Collaborate efficiently across teams and business units without ETL or integrations by sharing direct access to live data and Snowflake Native Apps.
- Discover the rich set of content within your organization, as well as data and apps available on Snowflake Marketplace, with a single Universal Search.

Security and Governance Details

Security is built into Snowflake. The platform is secure by design.

- Snowflake automatically encrypts all ingested data using AES-256 strong encryption. All communications are secured using TLS 1.2 and are controlled by network policies.
- Identity and access are managed through such enterprise features as SAML 2.0 based SSO, OAuth 2.0 delegated authorization, and SCIM for user management. For native Snowflake credentials, supported features include password policies, MFA, and key-pair authentication with key rotation.
- Robust role-based access control (RBAC) helps ensure that data and information can be accessed or operated on only by authorized users within an organization. Access control is applied to all database objects including tables, schemas, secure views, secure user-defined functions (secure UDFs), and virtual warehouses.
- Snowflake provides a comprehensive audit trail of all activities by all users.
- Governance features include automatic classification and tagging of sensitive data, dynamic data masking, including tag-based masking, row-access policies, access history, data lineage views and more.

SNOWFLAKE SECURITY & GOVERNANCE AT A GLANCE

1 Network Controls

- All communication secured using TLS 1.2 for all client communications, and controlled by Network Policies (IP Allowlisting)
- CSP Private Networking Integration
 - GCP Private Service Connect
 - AWS Privatelink
 - AWS VPC ID S3 policies
 - Azure Private Link
 - Azure subnet rules for Blob access

2 Identity & Access

- User Provisioning with SCIM
- Native Snowflake credentials
 - Multi-Factor Authentication
 - Key Pair Authentication
- Federated Identity
 - SAML 2.0-based SSO
 - OAuth 2.0 delegated authorization
- Password & Session Policies
- Service / Functional Accounts

3 Data Governance

- Built-in Features & partner integrations
- RBAC & DAC, Views & UDFs
- Column-Level Security
 - Dynamic data masking
 - External tokenization
- Row access policy
- Conditional policy
- Tagging
- Tag-based policy
- Classification
- Anonymization
- Object dependencies
- Access history

4 Encryption

- Customer data always encrypted in flight
- Data-at-rest always encrypted using a hierarchical key model
 - Rooted in the CSP's HSM
 - Automated key rotation & re-keying
 - BYOK with "Tri-Secret Secure"

5 Data Protection

- Account, region, cloud, and data-level recovery & failover
 - Fail-Safe
 - Time Travel
 - Cross-cloud & region replication & failover
 - AWS, Azure, GCP redundancy

6 Auditing

- Comprehensive audit trail for all activities by all users from login
 - Stored for 365 days in customer tamper proof area of account
- Access History for every object
- Pull or push models for export

7 Compliance & Legal



SOC 2 Type II
12 Month Coverage Period
SOC 1 Type II
6 Month Coverage Period



HITRUST Certified



Moderate (Available from OMB MAX)

- Full compliance report list
- Snowflake Security Addendum
- Snowflake's own Security Data Lake
- <https://www.snowflake.com/legal/> for DPA (GDPR), acceptable use, support, and more – about GDPR

For more details, visit <https://docs.snowflake.com/guides-overview-secure> and <https://docs.snowflake.com/guides-overview-govern>.

Snowflake AI Trust and Safety

Snowflake sets the industry standard for platform security, so customers don't have to lose sleep over security risks. All aspects of Snowflake's architecture, implementation, and operations are designed to

protect customer data in transit and at rest against both current and evolving security threats. Snowflake's approach to Generative AI is no different. Snowflake hosts and/or manages three types of large language models that power its AI features: its own proprietary LLMs, open-source LLMs, and licensed proprietary LLMs. Snowflake's AI features are subject to Snowflake's standard shared responsibility model for data protection, governance, and security. Snowflake does not use customer data to train any LLMs to be used across customers and does not use customer data to fine-tune any LLM except at the election or direction of the customer. Customer data never leaves the service boundary when using AI features. To learn more please visit <https://www.snowflake.com/en/legal/snowflake-ai-trust-and-safety/>.

CERTIFICATIONS AND COMPLIANCE

Snowflake is continuously expanding our [Security & Compliance Reports portfolio](#) as our customers request. The following is the current list of reports available to all customers and prospects under NDA:

SOC 2 Type II

The SOC2 Type 2 report is an independent auditor's attestation of the security controls that Snowflake has had in place during the report's coverage period. This report is provided for customers and prospects to review to ensure No Exceptions to the documented policies and procedures in the policy documentation.

SOC 1 Type II

The SOC1 Type 2 report, like the SOC2 Type 2 report, is an independent auditor's attestation of the financial controls that Snowflake has in place during the report's coverage period.

PCI-DSS

The Payment Card Industry Data Security Standards is a set of prescriptive requirements to which an organization must adhere in order to be considered compliant. Snowflake's Attestation of Compliance from our selected Qualified Security Assessor provides an independent auditor's assessment results after testing Snowflake's security controls.

HITRUST CSF

The Health Information Trust Alliance Common Security Framework (HITRUST CSF) serves to unify security controls based on aspects of US federal law (such as HIPAA and HITECH), certain state-specific laws and other industry-standard compliance frameworks into a single comprehensive set of baseline security and privacy controls, built specifically for healthcare needs. Snowflake participates in the HITRUST Shared Responsibility and Inheritance Program. With the Shared Responsibility Matrix (SRM), customers can now inherit Snowflake's HITRUST CSF certification provided that customers apply the controls detailed in the [HITRUST Alliance website](#). Customers should download the [Snowflake Custom HITRUST Shared Responsibility Matrix](#) to determine HITRUST controls that they can inherit as part of the shared responsibility model. Customers should refer to the [HITRUST webpage](#) for guidance on how to initiate an inheritance request.

CSA Star Level 1

[Cloud Security Alliance](#) (CSA) is a not-for-profit organization with a mission to "promote the use of

best practices for providing security assurance within Cloud Computing, and to provide education on the uses of Cloud Computing to help secure all other forms of computing.” Snowflake participates in the voluntary CSA Security, Trust & Assurance Registry (STAR) Self-Assessment to document our compliance with CSA-published best practices. The completed [CSA Consensus Assessments Initiative Questionnaire \(CAIQ\)](#) is found on the Cloud Security Alliance website.

ISO/IEC 27001, ISO 27017:2015 & ISO 27018:2019

The International Organization for Standardization provides requirements for establishing, implementing, maintaining, and continually improving an information security management system. Snowflake’s ISO Certificate is available for download by [clicking here](#). The statement of applicability additionally includes control objectives from the ISO 27017:2015 & ISO 27018:2019 framework.

ISO 9001:2015 Quality Management System

ISO 9001 standard sets forth universally acknowledged benchmarks for quality management, shaping the foundation for Snowflake to enhance its operations. Implementing ISO 9001 not only improves overall performance but also enables Snowflake to surpass customer expectations. By adhering to its stringent requirements, Snowflake showcases an unwavering commitment to quality. ISO 9001 guides the establishment, implementation, maintenance, and continuous improvement of a robust Quality Management System (QMS) at Snowflake, solidifying Snowflake’s commitment to delivering exceptional quality.

ISO/IEC 42001:2023

ISO/IEC 42001:2023 is the world’s first artificial intelligence management system standard. It provides guidance and requirements for the responsible management, development, and use of AI, along with criteria for the establishment, maintenance, and continuous improvement of an AI management system.

NIST SP 800-171

The National Institute of Standards and Technology (NIST) Special Publication (SP) 800-171 provides recommended security requirements for protecting the confidentiality of Controlled Unclassified Information (CUI) when the information is resident in nonfederal systems and organizations. Snowflake [documentation](#) details how Snowflake supports customers with NIST SP 800-171 compliance requirements.

SSDF

The Cybersecurity and Infrastructure Security Agency (CISA) Secure Software Development Framework ([SSDF](#)) reinforces secure by design principles advanced by CISA, Federal government partners, and international allies and requires software producers serving the federal government to confirm implementation of specific security practices.

Snowflake maintains service offerings that have completed a National Institute of Standards and Technology (NIST) Special Publication (SP) 800-218 SSDF assessment by a FedRAMP authorized third-party assessment organization (3PAO) with an accompanying attestation letter available upon request.

FedRAMP

The Federal Risk and Authorization Management Program, or [FedRAMP](#), is a government-wide

program that provides a standardized approach to security. Snowflake **SnowGov Regions** have achieved **FedRAMP Authorization**.

- **FedRAMP Moderate:** Federal Agencies may download Snowflake's FedRAMP Moderate Package from USDA Connect through community.connect.gov.
- **FedRAMP High:** Federal Agencies may download Snowflake's FedRAMP High Package by request via their Snowflake Account Executive.

Department of Defense (DoD) Impact Level 5 (IL5)

The Department of Defense (DOD) Cloud Computing Security Requirements Guide (SRG) outlines the security model and controls for the DOD's use of cloud computing. The U.S. military creates, stores, and operationalizes massive amounts of sensitive data. Protecting that data is a strategic priority and is the focus of the DOD Cloud Computing SRG framework. This framework is used to categorize information systems and data and to indicate the security requirements that data is subject to. Snowflake has received Provisional Authorization (PA) by the Defense Information Systems Agency (DISA) to operate at [Information Impact Level 5 \(IL5\)](#) on AWS GovCloud. This IL5 authorization allows Snowflake to offer authorized solutions to organizations requiring the highest levels of protection for Controlled Unclassified Information (CUI) within the DOD and related agencies. Agencies may download Snowflake's DoD package from eMASS.

GovRAMP (Moderate and High)

[GovRAMP](#) is a 501(c)6 nonprofit that standardizes security requirements for cloud offerings and verifies those cloud offerings for use by state and local governments and public education institutions through independent audits and continuous monitoring. The Snowflake offerings that are working towards or have achieved GovRAMP authorizations are included on the [Authorized Product List](#).

TX-RAMP

In the 87th Legislative Session, the Texas Legislature passed Senate Bill 475, requiring the Texas Department of Information Resources (DIR) to establish a state risk and authorization management program that provides "a standardized approach for security assessment, authorization, and continuous monitoring of cloud computing services that process the data of a state agency." To comply, DIR established a framework for collecting information about cloud services security posture and assessing responses for compliance with required controls and documentation. Texas Government Code 2054.0593 mandates that state agencies as defined by Texas Government Code 2054.003(13) must only enter or renew contracts to receive cloud computing services that comply with TX-RAMP requirements beginning January 1, 2022.

Texas Risk and Authorization Management Program (TX-RAMP) maintains an up-to-date inventory of cloud solutions and the authorization status of those solutions at the following location:

<https://dir.texas.gov/resource-library-item/tx-ramp-certified-cloud-products>.

GxP

GxP data integrity requirements (e.g.; 21 CFR 11) apply to life sciences organizations that produce regulated medical products including pharmaceuticals, medical devices, and mobile medical applications. Snowflake is **GxP compatible**, allowing life sciences customers to ensure data integrity and build GxP compliant solutions with the help of a secure, validated cloud data platform.

ITAR

International Traffic in Arms Regulations (ITAR) state that non-US persons are prohibited from physically or logically accessing the ITAR environment. A Third-Party Assessment Organization (3PAO) performed an audit to confirm that Snowflake's Microsoft Azure Government (MAG) and AWS GovCloud deployments provide an environment compliant with ITAR.

IRAP (Protected)

The Infosec Registered Assessors Program, or IRAP, is a program governed by the Australian Signals Directorate (ASD) of the Australian Government which endorses suitably-qualified cyber security professionals to provide relevant services which aim to secure broader industry and Australian Government systems and data. IRAP provides a security framework and an assessment methodology that enables Australian Government agencies and their customers to validate Snowflake's security control implementations and compliance against those requirements defined within the Australian Government Information Security Manual (ISM) developed by the Australian Signals Directorate (ASD). Snowflake employs IRAP assessors to validate Snowflake Australian systems effectiveness against the Information Security Manual at the Protected level.

CJIS

Snowflake's [SnowGov Regions](#) are ready and able to support customer compliance with the FBI's Criminal Justice Information Services (CJIS) Security Policy. The CJIS Security Policy provides federal and state agencies with a unified set of standards for the protection and safeguarding of Criminal Justice Information (CJI) in the cloud. Snowflake recognizes the importance of protecting CJI and works collaboratively with customers to satisfy CJIS requirements. Customers interested in learning more about how they can use Snowflake in connection with CJI can learn more here: [Criminal Justice Information at Snowflake](#).

IRS Publication 1075

Internal Revenue Service (IRS) Publication 1075 (IRS 1075) outlines the policies, practices, controls and safeguards to be employed by federal, state, and local agencies and contractors handling Federal Tax Information (FTI). Snowflake supports customer compliance with IRS 1075 in our FedRAMP-authorized SnowGov Regions. While there is no official certification for IRS 1075, Snowflake follows IRS Publication 1075 standards and works closely with our customers to meet the IRS's stringent regulatory requirements for the protection and safeguarding of FTI. For more information, please visit the [IRS Safeguards Program webpage on Cloud Computing](#).

TISAX (Assessment Level) AL 3

Developed by the ENX Association and published by the German Association of the Automotive Industry or VDA, Trusted Information Security Assessment Exchange or TISAX is a certification specifically designed to address the automotive industry's cybersecurity requirements. TISAX focuses on the secure processing of information from business partners, the protection of prototypes and data protection in accordance with the General Data Protection Regulation (GDPR) for potential business transactions between automobile manufacturers and their service providers or suppliers. TISAX was established in 2017 by VDA and the ENX Association. All organizations involved in business with major German automotive industry partners must obtain a TISAX certification. Assessment Level 3 is required for data with a very high need for protection, such as data classified as confidential or secret.

Snowflake's TISAX is scoped to Information Security and Assessment Level 3. For more information, please visit the [Official TISAX website](#).

CE+ (Cyber Essentials Plus)

Cyber Essentials Plus (CE+) is a United Kingdom government supported framework that helps protect organizations, regardless of size, against a wide range of the most common cyber attacks. The CE+ certification is required for organizations that plan to bid for central government contracts which involve handling sensitive or personal information or the provision of certain technical products and services. CE+ certification is supported by industry, including the Federation of Small Businesses, the Confederation of British Industry and a number of insurance organizations that offer incentives for CSPs holding this certification. CE+ provides the necessary technical controls and a related assurance framework conducted via an annual external assessment conducted by an accredited assessor. Achievement of the CE+ certification demonstrates Snowflake's commitment to mitigate the risk from common Internet-based threats and cyber security best practices. For more information, please visit the [Cyber Essentials website](#). You can view the current status of Snowflake's CE+ certification on the [IASME website](#).

C5 – Cloud Computing Compliance Controls Catalog

The Cloud Computing Compliance Controls Catalog (C5) is an audited standard establishing mandatory baselines for cloud security. The framework was created by the German Federal Office for Information Security (Bundesamt für Sicherheit in der Informationstechnik, or BSI). C5 was initially created for government agencies and organizations that work with the government to ensure that security baselines are met by their cloud service providers (CSPs). The private sector has also adopted this framework for evaluation of the security of their CSPs. The framework is based on ISO 27001, CSA, and BSI's IT-Grundschutz catalogs. The certification can be obtained for either the Basic requirements or Basic + Additional Criteria. Snowflake's C5 scope currently includes the Basic requirements.

K-FSI Cloud Service Provider Safety Evaluation and the Regulation on Supervision of Electronic Financial Transactions (RSEFT)

The Korean Financial Security Institute (K-FSI) performs the CSP Safety Evaluation in order to evaluate cloud service provider compliance with the Regulation on Supervision of Electronic Financial Transactions (RSEFT) regulation. They support the financial services industry in security assessments and assist in various areas that help create a secure environment for financial institutions. Snowflake's CSP Safety Evaluation is scoped to SaaS service controls. If customers will store or process unique private information (UPI) or protected credit information (PCI) on Snowflake or safety/reliability of electronic financial transactions are materially impacted by using Snowflake, then customers must review the CSP Safety Evaluation results and perform an analysis of vendor risk and submit documentation to the Financial Supervisory Service prior to utilizing Snowflake for these types of data.

WCAG 2.2AA

Snowflake has an event-driven process to update its Voluntary Product Accessibility Template (VPAT) rather than following a regular scheduled cycle. The company partners with Level Access, a specialized accessibility testing company, to conduct comprehensive VPAT assessments through approximately three weeks of rigorous testing.

Snowflake updates VPATs when the company makes significant platform changes, particularly major user interface updates. For example, Snowflake executed separate VPATs for its Classic UI and the New UI when it entered Public Preview, ensuring accessibility compliance across both interfaces.

While Snowflake conducts VPAT assessments periodically, the company maintains continuous accessibility improvements with weekly fixes as part of its ongoing accessibility commitment. Snowflake also updates VPATs in response to client requirements, particularly from federal government customers and other organizations that require Section 508 or WCAG compliance.

Snowflake manages the VPAT process cross-functionally, with its legal team and Product Management team serving as primary points of contact for client inquiries about accessibility compliance. This coordinated approach ensures Snowflake can respond quickly and accurately to client accessibility questions and requirements.

Snowflake Support aims to offer responsive, collaborative, and predictive support to help you maximize the Snowflake Data Cloud. The following are some examples our team can help troubleshoot:

- Product usage questions
- Feature enablement
- Troubleshooting failed queries
- Individual query performance analysis
- Supported third-party configuration
- Driver/connector support
- Liaison with engineering on product defects
- Product feedback
- Service alert notification

QUICKSTART PROGRAM FOR PUBLIC SECTOR

Snowflake offers Quick Start Packages for the Public Sector. Three offers, Quickstart, Quickstart+, and Accelerator, are available with fixed pricing.

- The Quickstart package serves as the basic entry-level starter offer, and includes:
 - Access to +30 on-demand Snowflake Training Modules
 - 1 seat in Snowflake Fundamentals Course
- The Quickstart+ package is a more enhanced, mid-tier offering for customers, and includes:
 - Access to +30 on-demand Snowflake Training Modules
 - 2 seats in Snowflake Fundamentals Course
 - QuickStart Accelerator
- The Accelerator is the most comprehensive and top-tier package available.
 - Access to +30 on-demand Snowflake Training Modules
 - 3 seats in Snowflake Fundamentals Course
 - QuickStart Accelerator

QuickStart Accelerator is a 2 week series of interactive workshops led by a Snowflake Professional Services Solutions Architect, designed to begin implementing a customer's primary use-case on Snowflake and build customer confidence.