

Brochure

The Modern DataOps Platform for the Snowflake Data Cloud

Unleash the Power of Data
Engineering With Spectra and
Snowflake Data Cloud

In Partnership With



The world's leading enterprises are moving to the cloud, and data is at the heart of this transformation. This shift is inevitable due to the growing push to adopt cloud-native technologies for operationalizing, securing, and monetizing data and analytics architectures.

The Snowflake Data Cloud

Snowflake Data Cloud is an end-to-end platform that solves the most challenging business problems with data. The Data Cloud is a hardened cloud service that scales easily and massively, allowing organizations to eliminate silos. An industry-leading performance, scalability, flexibility, efficiency, and ease-of-use have made Snowflake emerge as a platform of choice for companies looking to streamline their cloud adoption journey.

Enterprise DataOps for Speed and Scale

While selecting the right Data Cloud platform is critical, it is also imperative to combine it with a complete DataOps solution that provides an enterprise-ready toolset for a variety of workloads. A combined solution simplifies data pipeline and advanced analytics complexities, making it easier to achieve the desired benefits of data cloud transformation. This combination enables business stakeholders and data science experts to come together and drive real business value instead of wasting valuable cycles on building and maintaining custom data solutions.

Today, the focus is on executing enterprise-wide DataOps at speed and scale characterized by

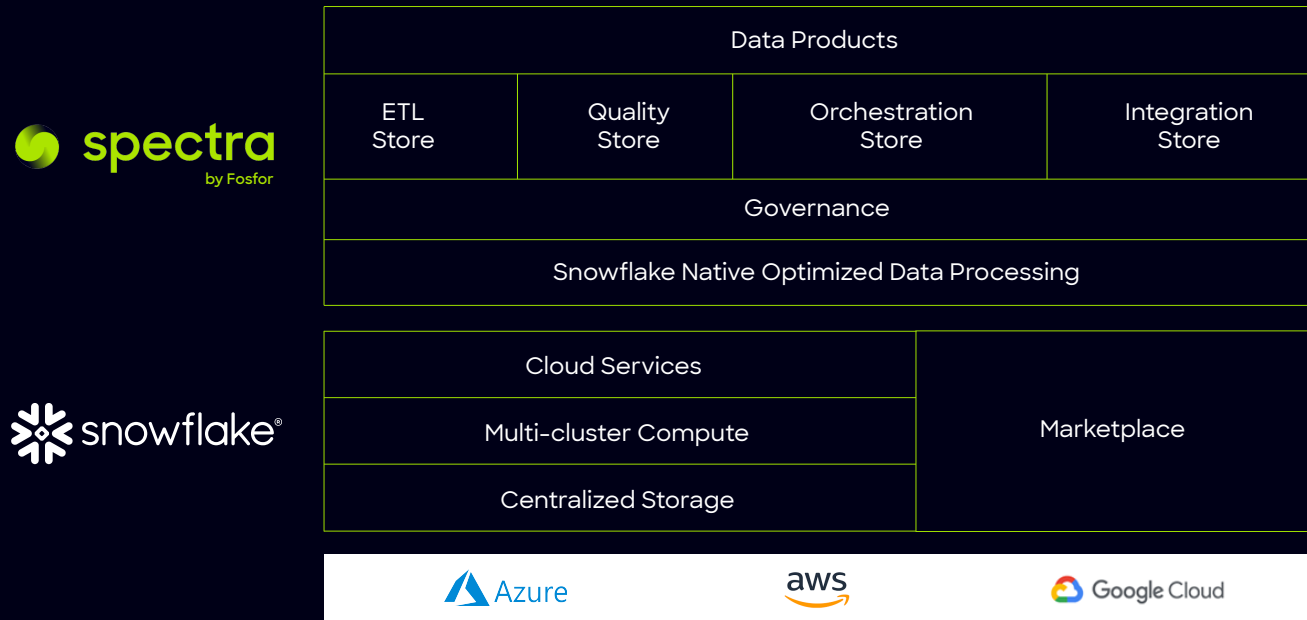
1. How end-users use data, rather than data ingestion and storage
2. Support for better data discovery and usage, rather than data extraction and loading alone
3. An ecosystem of data products available across the enterprise versus more limited data platforms
4. A domain-oriented approach focused on business needs of diverse and distributed teams rather than a processing-oriented approach focused on technical skills

Spectra + Snowflake Solution Architecture

Supercharge Enterprise DataOps on Cloud

Spectra, a comprehensive cloud-native DataOps platform, leverages advanced data processing engines like Snowflake to help data engineers and data analysts transform data assets to build data products by providing a self-serve, no-code platform.

The combination of Spectra and Snowflake Data Cloud enables enterprises to harness their data with industry-leading speed.



Spectra – Key Value Drivers

- Enables decentralized, domain-oriented teams to build data pipelines and data products independently
- Allows connections to many diverse sources of data without complex programming
- Facilitates extraction of insights from all types of data, including structured, semi-structured, and unstructured data
- Accommodates variable data lifecycles from creation, through usage, to eventual archive and deletion
- Allows for the centralization and governance of data in the cloud while still supporting decentralized DataOps
- Builds internal and external trust in data outcomes (e.g. data-driven decision making, building data applications etc.)
- Makes the delivery of data products easier for data scientists, analysts, and engineers

Spectra for Snowflake – Key Features

ETL Store

Drag-and-drop interface to orchestrate data pipelines

- **Data Connectors**
Data pipelining over 50 pre-built native connectors for cloud and on-premises source systems by enabling easy onboarding and automatic integration of these systems to Snowflake.
- **Data Pipelining**
Data processing functions including sort, lookup, join, transform, grouping, etc. are readily available.

Quality Store

Built-in AI-led data quality capabilities to increase trust in data

- **Data Profiling and Quality Checks**
Exploratory data analysis capabilities like data profiling, quality dashboards, and feature pipelines ensure trusted data on Snowflake. Reusable quality control can be configured across multiple domain-specific applications.

Orchestration Store

Workloads configured for Snowflake can be centrally scheduled and monitored

- **Scheduling and Monitoring**
This can be done centrally using a drag-and-drop canvas for data pipeline dependency and scheduling for workloads configured for Snowflake, while the execution of data pipelines can be monitored step-by-step or via error logs.

Integration Store

Seamless integrations with different B2B, AI / ML, and reporting applications

- **Data Consumption**
Transformed and processed data on Snowflake can be consumed by 3rd party applications for reporting AI / ML, ad-hoc analysis, and more, with inbuilt data-as-service capabilities through APIs and reverse ETL.
- **Custom Reusable Components**
Any business logic written in an external script (Python, Java, etc.) can be integrated into the data pipeline.

High-performance Hybrid and Multi-cloud Processing

Cloud-native data integration and transformations

- **Dynamic Execution**
Allows implementation of dynamic execution with optimized ELpT pipelines and a purpose-built dynamic execution framework.
- **Snowpark Execution**
Allows the creation of end-to-end data pipelines to perform data transformation natively within Snowflake by leveraging Snowpark capabilities.

Data Products

Decentralized, domain-oriented teams building data pipelines and data products

- **Domain Data Pipelines**
Leverage pre-built domain data pipelines to accelerate the data product journey for different domains such as manufacturing, insurance, pharma, CPG, etc.
- **Data Exploration and Discovery**
The ability to independently maintain and automatically version control data products simultaneously allow strong governance, access management, and discovery of data across different domains and applications.

Governance

Integration with enterprise governance and security measures

- **Authentication and Authorization**
Platform access is based on user personas, ensuring effective user management to manage and sync users, groups, roles, and actions from the enterprise LDAP or AD; the platform allows role-based access to users, including developer, operator, and administrator roles.
- **Lineage**
Data lineage is maintained from source to target with related transformation information across applications and processes.

Driving Value Together



- **Fast-track Snowflake Adoption**
Cloud-native integration support for Snowflake on AWS, Azure, and GCP
- **Expedite Time-to-Market Data Products**
Pre-built domain pipelines for operationalizing data products
- **Reusability**
Centralized QC and domain pipelines aid reusability across all Snowflake apps
- **Speed and Performance**
Cloud-native Push-Down Optimization (PDO) offloads processing in the Snowflake engine
- **Compute and Storage Resource Scalability**
Empower instant scale up/down to manage millions of records without latency

To Learn More Contact: info@fosfor.com

The Fosfor Product Suite is the only end-to-end suite for optimizing all aspects of the data-to-decisions lifecycle. Fosfor helps you make better decisions, ensuring you have the right data in more hands in the fastest time possible. The Fosfor Product Suite is made up of Spectra, a comprehensive DataOps platform; Optic, a data fabric to facilitate data discovery-to-consumption journeys; Refract, a Data Science and MLOps platform; Aspect, a no-code unstructured data processing platform; and Lumin, an augmented analytics platform. Taken together, the Fosfor suite helps businesses discover the hidden value in their data. The Fosfor Data Products Unit is part of LTI, a global technology consulting and digital solutions company with hundreds of clients and operations in 31 countries. For more information, visit [Fosfor.com](https://fosfor.com).