

Overview

Managed Data Services for Snowflake



Challenges managing your data platform

Typically, Snowflake based data platform involves multiple ecosystem tools for data integration, workflow management, BI, data science, data security, and governance.

Customers can risk their data quality, data freshness, high availability and data security / access control over the data if the platform is not managed well.

The management of the entire data platform is complex.

Persistent addresses your challenges

Support multiple tools for observability, processing, business intelligence, visualizations, AI and ML, and health monitoring and security governance for your data platform.

For a monthly cost, we provide a managed service to mitigate issues of cost, complexity and execution.

We offer an accelerated process — to get the managed service operational within 6 weeks.

Persistent is a leader in the Snowflake Ecosystem

Prebuilt Accelerators for Operations Management

Data Quality Measurement

Cost Monitoring Dashboard, Reports and Predictive Algorithms

Integration with Open Source Tools for Infrastructure / Services Monitoring

3rd Party Tool Integration and Expertise for Data Security and Access Control Management and Monitoring

Proven Expertise with 20+ Active Snowflake Engagements

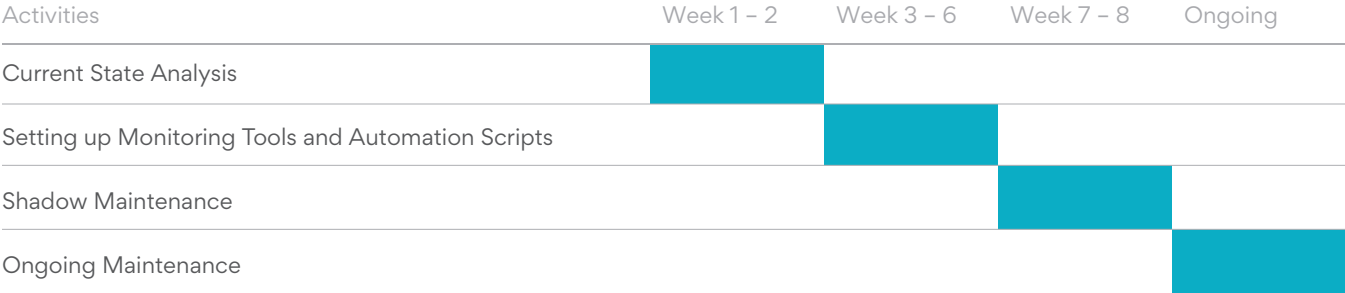
Deep Knowledge of Snowflake Connectors — Persistent runs the Snowflake Ready Technology Validation Program

What we offer — Managed data platform packages for Snowflake

DataOps Category	Standard	Enterprise	Premium	Comments
Infrastructure Creation	Y	Y	Y	
Infrastructure Management	Y	Y	Y	
Data Pipeline Monitoring	Y	Y	Y	
Data Quality Monitoring	Y	Y	Y	
Data Catalog and Observability Monitoring	Y	Y	Y	
Cost Monitoring	Y	Y	Y	
Data Access Monitoring	Y	Y	Y	
Number of Snowflake Users	up to 40	up to 100	> 100	
Number of Data Pipelines	up to 200	up to 500	> 500	Airflow and Matillion DAGs
Support	8x5	8x7	24x7	

Pricing starts at \$195,000 annual cost. Pricing for data pipeline development, data quality rule configuration, and catalog implementation is additional. Cloud usage is charged at actuals. Talk to our sales team for additional details.

Methodology



Snowflake DataOps

Category	Task	Type	Description
Infrastructure Creation	Creation of new environment (Dev / Stage / Prod)	Support ticket	New environment creation with specific tools (e.g. S3, Airflow, Snowflake, Jupyter, Matillion)
	Patching / Upgrades for deployed softwares	Ongoing maintenance	Deploying patches, etc., on all environments
	Tool integrations, library deployments	Support ticket	
	Resource sizing	Support ticket	Helping with resource sizing (Snowflake Virtual Warehouse configuration, Matillion configuration etc.)
	SSO Integration	Ongoing maintenance	
	Design Access Control (RBAC / Authentication) framework	Support ticket	
	Scheduling backups	Support ticket	
	Putting usage / cost alerts	Support ticket	Cost alerts for the entire data domain infrastructure
Infrastructures Operations	Adding users to environments and tools	Support ticket	
	Creating roles in various tools	Support ticket	IAM roles, Snowflake roles etc.
	Implementing Access Control (new access, restrict access etc.)	Support ticket	
	Analyzing usage reports for security risks	Ongoing maintenance	e.g. unnecessary access, failed login attempts, login anomalies etc.
	Analyzing resource sizing	Ongoing maintenance	e.g. underutilized Snowflake Virtual Warehouse, overutilized / underutilized Matillion instances etc.
	Monitoring storage	Ongoing maintenance	e.g. Snowflake internal / external stage, its unusual growth etc.
	Monitoring slow and long running queries	Ongoing maintenance	

Category	Task	Type	Description
Data Pipeline Monitoring	Metrics	Ongoing maintenance	e.g. Airflow will be configured to send various counters which will be monitored. The counters include #StartedJobs, #CompletedJobs, #HeartBeatFailures, #ZombiesKilled, #DagRefreshErrors etc. For more details, refer to: — https://airflow.apache.org/docs/apache-airflow/stable/logging-monitoring/metrics.html
	Logs	Ongoing maintenance	e.g. watching Airflow logs (generated by web servers, schedulers and workers) for errors, warnings
	Schedules	Ongoing maintenance	Monitoring jobs schedules for ontime start, completion and intermediate step completion
	Dependencies	Ongoing maintenance	Monitoring entire workflow in the pipeline for dependent tasks completion
Data Quality Monitoring	Metrics	Ongoing maintenance	e.g. Monitoring %objects under DQ, %rows failed, %rules failed, first time rule failure, consistent rule failure etc., metrics for each dataset under DQ
	Logs	Ongoing maintenance	Watching DQ data pipeline execution logs
	Schedules	Ongoing maintenance	Monitoring execution of DQ pipeline
	Trend	Ongoing maintenance	Monitoring the overall trend for DQ score on each dataset and alerting for outliers / degrades
Data Catalog and Observability Monitoring	Catalog object metrics	Ongoing maintenance	<ol style="list-style-type: none"> 1\ Number of objects in the catalog and its trend 2\ Number of fields in each object, number of NULL and not NULL fields
	Catalog usage metrics	Ongoing maintenance	<ol style="list-style-type: none"> 1\ Number of consumers of the catalog
Cost Monitoring	Cost monitoring dashboards	Ongoing maintenance	Cost Monitoring Dashboards
	Cost alerts	Ongoing maintenance	e.g. detecting sudden surge, detecting outliers in cost

Category	Task	Type	Description
Data Access Monitoring	User login activity	Ongoing maintenance	Sudden surge in the user logins, multiple failed login attempts
	Client tools connectivity	Ongoing maintenance	Login from unapproved enterprise tools (e.g. login from GoLang driver if it is not approved)
	Data load / unload activity monitoring	Ongoing maintenance	Watching the data unload activities for any outliers
	Data sharing	Ongoing maintenance	Monitoring of data sharing activities and users with data sharing privileges

Unlock the value of your data with our Snowflake Managed Services

[Learn More](#)

About Persistent

We are a trusted Digital Engineering and Enterprise Modernization partner, combining deep technical expertise and industry experience to help our clients anticipate what's next. Our offerings and proven solutions create a unique competitive advantage for our clients by giving them the power to see beyond and rise above. We work with many industry-leading organizations world-wide including 14 of the 30 most innovative US companies, 80% of the largest banks in the US and India, and numerous innovators across the healthcare ecosystem. Our company fosters a values-driven and people-centric work environment. Our strength of over 21,500 employees is spread over 18 different countries across the globe.

USA

Persistent Systems, Inc.
 2055 Laurelwood Road, Suite 210
 Santa Clara, CA 95054
 Tel: +1 (408) 216 7010
 Fax: +1 (408) 451 9177
 Email: info@persistent.com

India

Persistent Systems Limited
 Bhageerath, 402
 Senapati Bapat Road
 Pune 411016
 Tel: +91 (20) 6703 0000
 Fax: +91 (20) 6703 0008



Persistent