RapidCanvas

PARTNER CASE STUDY | RAPIDCANVAS + SNOWFLAKE

From Governed Snowflake Data To Shipped Agents

How RapidCanvas reads Snowflake in place, pushes compute into the customer's own warehouse, and turns governed tables into scored pipelines, hosted APIs and applications teams use.

2	1	September 2026
use cases	in production	edition

Real AI transformation, accelerated.



AT A GLANCE

RapidCanvas + Snowflake

Enterprises already trust Snowflake for governance, lineage and access control. RapidCanvas does not rebuild it. Our Hybrid Approach™, an agentic platform paired with expert-in-the-loop delivery, reads Snowflake in place through a native connector, pushes compute into the customer's own virtual warehouse, and turns governed schemas into scored pipelines, hosted data APIs and applications operators use.

The two stories here show both directions of that relationship. A RapidCanvas-built prospect-scoring pipeline writes ranked results back into Snowflake with transactional merge semantics and key-pair authentication. Separately, Snowflake is a first-class RapidCanvas data source available to every project, with Arrow batch reads and backup-then-restore writes that never leave a table empty. Snowflake stays the system of record; RapidCanvas ships the delivery layer on top.

2

Snowflake use cases documented

10

platform data source types, including Snowflake

3

Snowflake write modes: append, truncate, merge

PARTNERSHIP

RapidCanvas engages with Snowflake customers directly today; Snowflake Cloud Partner status is on the roadmap, not yet in place.

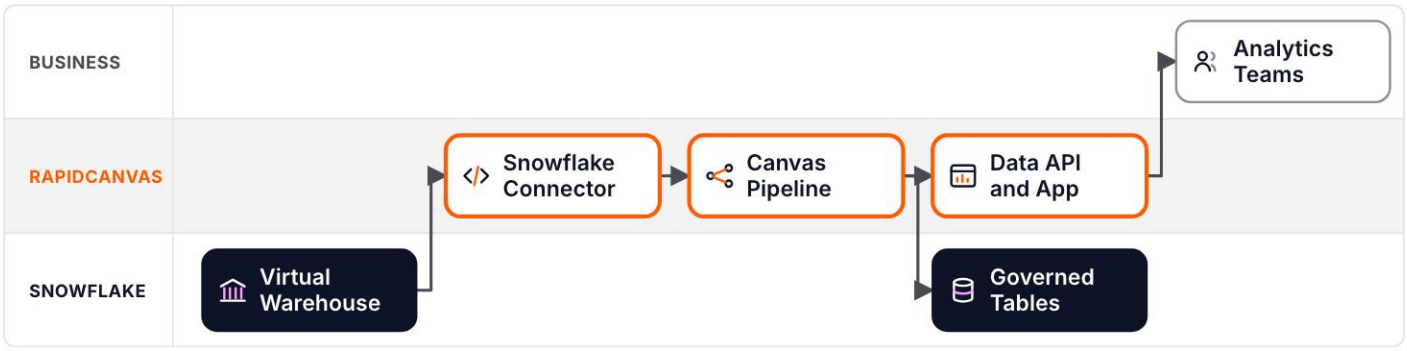
When RapidCanvas Is a Fit

1. Snowflake holds governed data but no team to turn it into scored, agent-ready applications.
2. The customer needs write-back into Snowflake with transactional merge semantics, not a one-way export.
3. Multiple warehouses or teams need one connector pattern instead of a bespoke client per project.
4. The account runs Snowflake and wants agents shipped without moving data out of the warehouse.

REFERENCE ARCHITECTURE

How RapidCanvas Builds on Snowflake

Snowflake stays the governed source of truth; RapidCanvas reads it through the native connector, pushes compute down, and ships the pipelines, APIs and applications on top.



Source Of Truth Stays Put

Snowflake keeps governance, lineage and access control; RapidCanvas reads through the standard connector rather than copying data out.

Compute Runs In Snowflake

Parameterised SQL executes in the customer's virtual warehouse; RapidCanvas receives result sets, not raw tables, so warehouse spend and governance stay with the customer.

Credentials Never Touch Code

Role-scoped Snowflake users are held in RapidCanvas Secrets and decrypted only when a connection opens, with key-pair authentication preferred over passwords.

Use Cases in This Document

NO.	USE CASE	INDUSTRY	STATUS
01	Scored Prospects Written Back Into Snowflake Automatically	Electronics manufacturing, global	**In production**
02	Snowflake Becomes A Native RapidCanvas Data Source	All industries, global platform capability	**Reference pattern**

USE CASE 01 | ELECTRONICS MANUFACTURING, GLOBAL ● IN PRODUCTION

Scored Prospects Written Back Into Snowflake Automatically

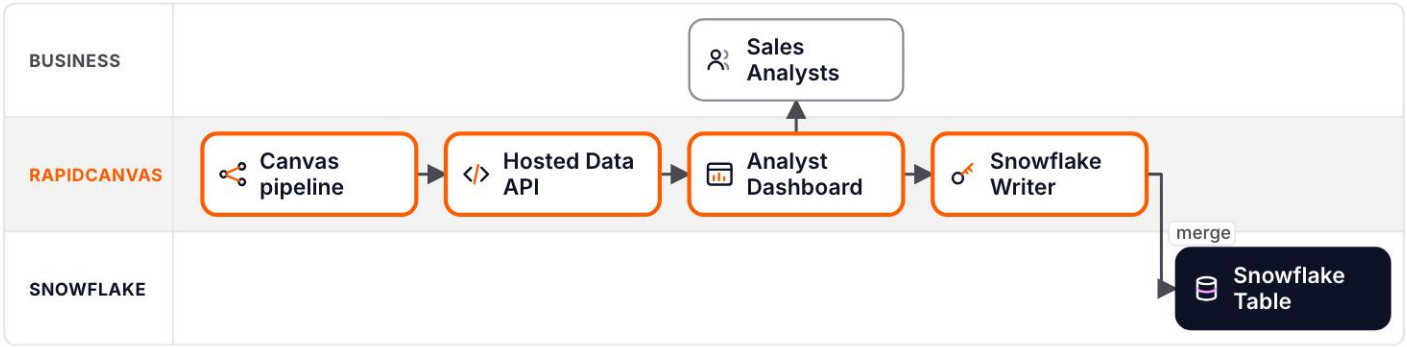
Snowflake Connector for Python | Snowflake key-pair authentication | Snowflake merge and truncate writes | Canvas pipelines | Data Apps | Connectors

The Challenge

An electronics manufacturer had assembled roughly ten thousand raw prospect companies, but no ranked, evidence-backed list. Sellers could not separate fit from intent or estimate deal size, and feedback on scores got lost. The customer's analytics estate runs on Snowflake, so scored leads had to land there.

Solution Implemented

RapidCanvas built a weekly Canvas orchestrator that scores fit and intent, predicts order value with a regression model, clusters companies, and promotes results into a history table with slowly changing dimensions. A hosted data API serves the list to a dashboard with an evidence drawer. When an analyst pushes results, a Snowflake write client authenticates with an encrypted key and writes in append, truncate or merge mode.



How It Works

1. A weekly Canvas orchestrator scores fit and intent, and predicts order value with a regression model.
2. Results are clustered, ranked, and promoted into a history table with slowly changing dimensions.
3. A hosted data API serves the ranked list to a dashboard with an evidence drawer.
4. An analyst-triggered push authenticates with an encrypted key and writes to Snowflake in merge mode.

Results and Benefits

- The pipeline scores about ten thousand prospect companies in a single weekly run.
- Scored, ranked leads write directly into Snowflake with transactional merge semantics, not a manual export.
- Analysts get evidence-backed fit and intent scores instead of a flat, unranked list.

HOW IT RUNS ON SNOWFLAKE

Snowflake runs the virtual warehouse and stores every scored table, so consumption grows with each pipeline run and analyst push.

USE CASE 02 | ALL INDUSTRIES, GLOBAL PLATFORM CAPABILITY ● REFERENCE PATTERN

Snowflake Becomes A Native RapidCanvas Data Source

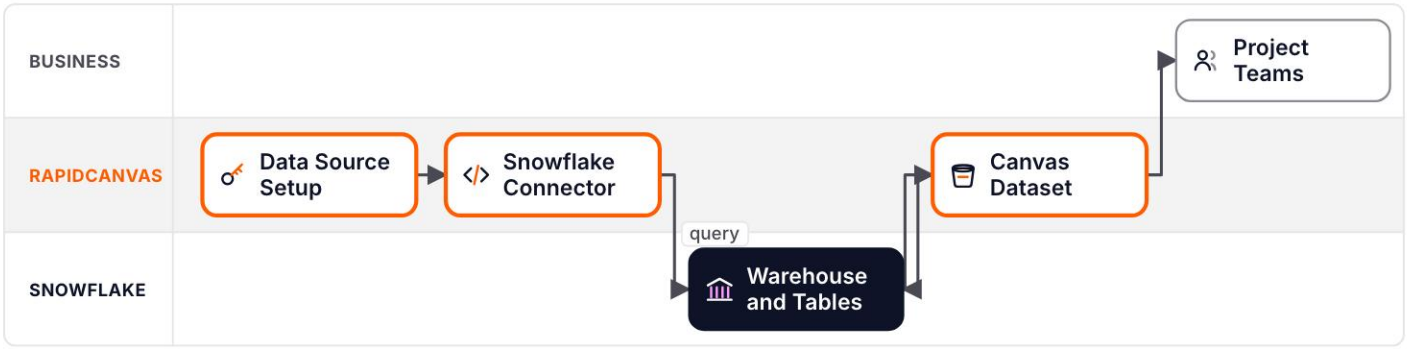
Snowflake Connector for Python | Snowflake SQLAlchemy | Snowflake key-pair authentication | Snowflake Arrow batch reads | Connectors | Canvas pipelines

The Challenge

Project teams kept building a Snowflake client for every solution, duplicating authentication and safe write logic each time. The platform needed one connector with the same experience as its other data sources, safe defaults for large reads and writes, and credentials that never sit in code.

Solution Implemented

RapidCanvas built Snowflake into the platform as a native data-source type on the official connector and Snowflake SQLAlchemy. A user configures account, warehouse, database, schema and role once; options are stored encrypted, with key-pair authentication preferred over passwords. Full reads stream Arrow batches into memory-bounded parquet chunks. Writes append, replace or truncate, backing up the table first and restoring it if anything fails.



How It Works

1. A user configures account, warehouse, database, schema and role once, with credentials stored encrypted.
2. Every connection is tagged with a RapidCanvas identifier so administrators can attribute the workload.
3. Full reads stream Arrow batches into memory-bounded parquet chunks that feed straight into Canvas pipelines.
4. Truncate writes back up the table first and restore automatically if the write fails midway.

Results and Benefits

- Snowflake is one of ten data-source types available to every RapidCanvas project today.
- Every project shares one tested connector instead of a bespoke client per solution.
- Writes never leave a target table empty, even when a truncate fails partway through.

HOW IT RUNS ON SNOWFLAKE

Every RapidCanvas project on Snowflake drives virtual warehouse compute and keeps governed tables as the destination for pipeline output.

WORKING TOGETHER

RapidCanvas on Snowflake

Every engagement runs on the Hybrid Approach™: the RapidCanvas agentic AI platform builds the pipelines, agents and Data Apps, and RapidCanvas experts work alongside the customer's team on data modeling, validation and the last mile into production.

RAPIDCANVAS BENCHMARKS ACROSS ENGAGEMENTS

4-8 weeks
to first production value

10x
faster time to value than a traditional custom build

Up to 80%
lower cost than a traditional custom build

Connect Natively

Snowflake is a supported RapidCanvas data-source type on the official Python connector, with account, warehouse, database, schema and role as connection options; datasets import from tables or queries and pipeline outputs export back.

Run In Any Cloud

RapidCanvas runs as SaaS or privately on AWS, Azure or Google Cloud, sitting alongside the customer's Snowflake account rather than requiring data to move anywhere first.

Keep Governance In Place

Snowflake role-based permissions govern every RapidCanvas query, and each connection carries a RapidCanvas application tag so administrators can see exactly what ran.

Work With Snowflake Directly

The commercial relationship with Snowflake is direct today, project by project; Snowflake Cloud Partner status is on the roadmap and not yet in place.

Snowflake Services in These Stories

Snowflake virtual warehouses
Runs parameterised SQL pushed down from RapidCanvas pipelines

Snowflake role-based access control
Scopes every RapidCanvas credential to a specific role

Snowflake application tagging
Tags every RapidCanvas connection so admins can attribute usage

Snowflake session parameters
Normalizes timestamp handling for every RapidCanvas query

Snowflake Connector for Python
Official client library behind the RapidCanvas connector

Snowflake SQLAlchemy
Backs chunked dataset builds alongside the connector

Snowflake key-pair authentication
Encrypted PEM credentials held in RapidCanvas Secrets

Snowflake Arrow batch reads
Streams result sets in sixteen-thousand-row chunks

Snowflake merge and truncate writes
Writes scores back with transactional, never-empty semantics

Customer names are withheld. Each story describes a real RapidCanvas deployment, verified against the delivered code. Delivery status is shown as of September 2026.

RapidCanvas

NEXT STEP

See How RapidCanvas Works on Snowflake

See how RapidCanvas works with Snowflake, or book a discovery call to assess fit for your data.

Book a Discovery Call

rapidcanvas.ai



Real AI transformation, accelerated.