



Automating Database Operations with Cortex Code

> CONF42

TIMESTAMP: MAY 2026



Safe Harbor and Disclaimers

Other than statements of historical fact, all statements contained in these materials and any accompanying oral commentary (collectively, the “Materials”) are forward-looking statements within the meaning of Section 27A of the Securities Act of 1933, as amended, and Section 21E of the Securities Exchange Act of 1934, as amended, including statements regarding (i) Snowflake’s business strategy, plans, opportunities, or priorities; (ii) Snowflake’s new or enhanced products, services, and technology offerings, including those that are under development or not generally available; (iii) market size and growth, trends, and competitive considerations; (iv) Snowflake’s vision, strategy, and expected benefits relating to artificial intelligence (“AI”), Snowpark, Snowflake Marketplace, the AI Data Cloud, and AI Data Clouds for specific industries or product categories, including the expected benefits and network effects of the AI Data Cloud; and (v) the integration, interoperability, and availability of Snowflake’s products, services, or technology offerings with or on third-party platforms or products, including public cloud platforms. These forward-looking statements are subject to a number of known and unknown risks, uncertainties, and assumptions, including those described under the heading “Risk Factors” and elsewhere in the Annual Reports on Form 10-K and the Quarterly Reports on Form 10-Q that Snowflake files with the Securities and Exchange Commission from time to time. In light of these risks, uncertainties, and assumptions, the future events and trends discussed in the Materials may not occur, and actual results could differ materially and adversely from those expressed or implied in the forward-looking statements. As a result, you should not rely on any forwarding-looking statements as predictions of future events. Forward-looking statements speak only as of the date the statements are first made and are based on information available to Snowflake at the time those statements are made and/or management’s good faith belief as of that time. Except as required by law, we undertake no obligation, and do not intend, to update the forward-looking statements in these Materials.

Any future product or roadmap information (collectively, the “Roadmap”) is intended to outline general product direction. The Roadmap is not a commitment, promise, or legal obligation for Snowflake to deliver any future products, features, or functionality; and is not intended to be, and shall not be deemed to be, incorporated into any contract. The actual timing of any product, feature, or functionality that is ultimately made available may be different from what is presented in the Roadmap. The Roadmap information should not be used when making a purchasing decision. In case of conflict between the information contained in the Materials and official Snowflake documentation, official Snowflake documentation should take precedence over these Materials. Further, note that Snowflake has made no determination as to whether separate fees will be charged for any future products, features, and/or functionality which may ultimately be made available. Snowflake may, in its own discretion, choose to charge separate fees for the delivery of any future products, features, and/or functionality which are ultimately made available.

The Materials may contain information provided by third-parties. Snowflake has not independently verified this information, and usage of this information does not mean or imply that Snowflake has adopted this information as its own or independently verified its accuracy.



NOTICE

This talk is solely the personal opinion, tips and ideas of Timothy Spann and does not represent the views or opinions of Snowflake.



Tim Spann

paasdev.bsky.social
@PaasDev // Blog: datainmotion.dev

Senior Solutions Engineer, Snowflake
NY/NJ/Philly - Cloud Data + AI Meetups

PAST EXPERIENCE:

ex-Zilliz ex-Pivotal ex-Cloudera ex-HPE

ex-StreamNative ex-EY ex-Hortonworks

<https://medium.com/@tspann>

<https://github.com/tspannhw>



 **DZone.** REFCARDS TREND REPORTS E

Top IoT Experts

 **Tim Spann**
Principal Developer Advocate,
Cloudera

<https://github.com/tspannher/SpeakerProfile>
Tim Spann is a Principal Developer Advocate
in Data in Motion on Cloudera. He works
with Apache HPL, Apache Pulsar, Apache...



AI + Streaming Weekly by Tim Spann



snowflake

<https://bit.ly/32dAJft>

This week in Apache NiFi, Apache Polaris, Apache Flink, Apache Kafka, ML, AI, Streamlit, Jupyter, Apache Iceberg, Python, Java, LLM, GenAI, Snowflake, Unstructured Data and Open Source friends.

Snowflake Cortex AI is Easy, Connected, Trusted



SNOWFLAKE CORTEX AI

AGENTIC BUSINESS INSIGHT

Agent Apps
Snowflake Intelligence PU Soon

Agent APIs
Cortex Agents GA Soon

STATE OF THE ART RETRIEVAL

Unstructured Data Retrieval
Cortex Search

Structured Data Retrieval
Cortex Analyst

SCALABLE AI PROCESSING

Multimodal AI-powered SQL
Cortex AISQL PU

Document Processing
Document AI, Parse, Embed

MODELS

OpenAI

ANTHROPIC

Meta

MISTRAL AI

deepseek

snowflake®

GOVERNANCE

RBAC

Guardrails

Evaluations

Monitoring

AI Gateway

DATA



DATA INGESTION

Fully managed extraction,
load, preprocess, and
activation of multimodal data



© 2022



Agenda

📺 What is Cortex Code?

⚙️ Architecture & Capabilities

⚡ Live Demo: Schema → dbt → Agents → CI/CD

📁 Key Takeaways & Resources

> **INITIALIZING: Automating Database Operations with Cortex Code CLI**



Extensible & Built on Open Standards



AGENTS.md



MCP



**AGENT
SKILLS**

CUSTOMIZE

CONNECT

EXTEND



The core pillars of high-impact AI organizations



AI in Every Workflow

*Unlock intelligence in the workflows of **every team** and **every agent**.*



Business Logic and Context

Ensure AI decisions align with your goals, rules, and company context.



Unified Data Foundation

Unify all data into one trusted foundation, regardless of where it lives.



One platform to power your entire AI strategy

All Teams



All Tools



All Data



AI DATA CLOUD

Unified Platform and Connected Ecosystem

Data Engineering

Analytics

AI/ML

Applications & Collaboration

Fully Managed | Cross-Cloud | Interoperable | Secure | Governed

Predictable Performance at Scale



Faster Time to Market



Enterprise-Grade Reliability



Business logic

and

context

Snowflake Cortex AI

Governed LLMs

ANTHROPIC

Mistral AI

OpenAI

Meta

Gemini

Custom & Open

Business Context

Business Logic

BI & Analytics

Agentic Experiences


ON-DEMAND
COMPUTE ENGINE





Unified data foundation

Eliminate data silos to scale intelligence and drive data-driven decisions anywhere



For Data Teams on Snowflake, The Right Tool ...

UNDERSTANDS SNOWFLAKE



Has the latest knowledge of your specific data catalog; databases, schemas, tables, semantic models and more

INTELLIGENCE

IS CONTEXT AWARE



Knows all the data you're working with, retains your conversational history, and where you are in the UI as you navigate

RELEVANCE

AUTOMATES COMPLEX TASKS



Intelligently optimizes your requests to solve real jobs, from building an application to managing cost or governing data

EFFICIENCY

IS SECURE BY DESIGN



Adheres to Snowflake RBAC, ensuring it only sees and acts on the data and objects you have access to

GOVERNANCE



Introducing Cortex Code

Snowflake-Native AI Coding Agent for Data Teams



Accelerates End-to-End Development

Streamline the entire data lifecycle by using natural language to automate data engineering, machine learning, and agent building tasks



Understands your Enterprise Data Context

Deeply aware of Snowflake's data, compute, governance, and operational semantics empowering everyone from data dabblers to domain experts



Is Open, Extensible, and Enterprise-Ready

Operates on secure Snowflake data and leverages RBAC, natively supports MCP, agents.md, and Agent Skills with access to frontier models like Claude Opus 4.6 & Sonnet 4.6, and GPT 5.2

SNOWSIGHT



Build a fraud detection ML model to help me predict credit card fraud in real-time



CLI

```
> Create dbt mart joining customers/orders for LTV, including schema.yml with unique and not-null tests
```



What Can You Do with Cortex Code?



CORTEX CODE

ADMIN



Data Discovery

Search the catalog for data, with or without semantic views.



Admin Tasks

Manage catalog, permissions, users, costs, governance etc.

BUILD ON SNOWFLAKE



Build AI Agents

Generate realistic synthetic data and create backend logic for agents.



Analytics & Apps

Streamlit, Notebook, Data engineering and pipelines.

DATA ENGINEERING



Create Manage & Optimize

Build dbt projects, pipelines and optimize existing projects.



Code Migration & Refactoring

Convert code/pipelines from other platforms to Snowflake.

...AND MUCH MUCH MORE!



BUILDING YOUR OWN AGENTIC AI DATA PIPELINE



Cortex Code: Snowflake-native AI coding agent

Core Capabilities



SQL
Execution



Snowflake
Native



Streamlit



Python/Snowpark



Multi-Account

Integrations



MCP
(Model Context
Protocol)



Documentation
Access




External APIs

Agent Features



Skills (Custom & Built-in)




Planning Mode

Custom Instructions
(agents.md)



Dev Tools



Terminal/Shell
(/sh)



Local File
Access



Git & dbt
Support

NEURAL INTERFACE TERMINAL

> INITIALIZING AGENTIC PIPELINE...

[STATUS: ONLINE]

[CORTEX_CODE: ACTIVE]

[AUTH: SNOWFLAKE_RBAC]

> SYSTEM READY_

SCANNING SCHEMAS
[98.2% COMPLETE]

LOAD_SOURCE_01: SUCCESS
CLI_CONNECTION: SUCCESS
DBT_MODELS: COMPLETED
PYTHON_EXEC: PENDING



WHAT IS CORTEX CODE?



Intelligent AI Agent

AI-driven agent purpose-built for Snowflake environments.



Native Understanding

Deep knowledge of Snowflake RBAC, schemas, and architecture best practices.



Unified Access

Seamlessly available within Snowsight UI or via high-performance CLI tools.



Multi-Language Support

Native execution and optimization for SQL, Python, dbt, and beyond.

CORTEX CODE CLI – KEY FEATURES



Snowflake Integration

Direct account access for seamless data operations.



Local File Access

Read and write directly to your local repositories.



Tool Orchestration

Automated Bash, Git, and SQL execution workflows.



Agent Customization

Defined via AGENTS.md, Skills, and Lifecycle Hooks.



Enterprise-Grade Security

Hardened via RBAC, secure sandboxing, and human-in-the-loop approval systems.

ARCHITECTURE OVERVIEW



Natural Language Loop

Continuous feedback loop from natural language input to AI agent execution.



40+ Built-in Skills

Pre-configured domain expertise for rapid deployment and deep specialization.



MCP Integration

Model Context Protocol (MCP) for seamless external tool and data connectivity.



Parallel Subagents

Autonomous execution through subagents for complex, parallel task handling.



Lifecycle Customization

Advanced hooks for deep customization of the agent's operational lifecycle and behavior.

INSTALLATION

 macOS / Linux:

```
curl -LsS  
https://ai.snowflake.com/static/  
cc-scripts/install.sh | sh
```

 Windows:

```
irm  
https://ai.snowflake.com/static/  
cc-scripts/install.ps1 | iex
```

 Installs to: ~/.local/bin/cortex



Connecting to Snowflake



Uses `~/.snowflake/connections.toml`



Compatible with Snowflake CLI connections



Setup wizard on first launch



Multiple connections supported



First Commands

-  `cortex` — Start interactive REPL
-  `cortex -p "query"` — Non-interactive mode
-  `cortex --resume last` — Resume session
-  `cortex --connection <name>` — Specific connection



Natural Language → DDL

- “Create a database called DEMO_AUTOMATION_DB”

- Cortex Code generates:

```
CREATE DATABASE IF NOT EXISTS ...
```

- Validates your role has proper privileges

- Executes and confirms

Creating Schemas

- “Create three schemas: RAW, STAGING, ANALYTICS”
- Understands medallion architecture pattern
- Generates three CREATE SCHEMA statements
- Applies best practices automatically



Creating Tables with Context

>_ "Create a customers table with e-commerce fields"

 Infers data types from business context

 Adds constraints and defaults

 Maintains relationships between tables

>_



Schema Verification

“Show me the DDL for all tables”

Uses GET_DDL() function

Provides documentation baseline

Ready for version control





Connecting to Snowflake



Uses ~/.snowflake/connections.toml



Compatible with Snowflake CLI connections



Setup wizard on first launch



Multiple connections supported



Scaffolding a dbt Project



“Create a new dbt project called demo_analytics”

Generates `dbt_project.yml`, `profiles.yml`



Creates `/models` directory structure

Sources, staging, and mart layers



Deploying to Snowflake

```
> INITIALIZING DEPLOYMENT...
```

```
$ snow dbt deploy demo_analytics
```


```
[✓] Creates DBT PROJECT object
```

```
[✓] Versioned (rollback capable)
```

```
[✓] EXECUTE DBT PROJECT to run
```

```
[✓] Monitor in Snowsight
```

```
STATUS: STABLE_V2.0
```



Complex SQL Generation

input



“Write a customer cohort retention analysis”



Automates complex SQL construction from natural language inputs.



Outputs production-ready, modular SQL utilizing **Common Table Expressions (CTEs)**.



Accurately handles intricate window functions and time-series date manipulation.



Ensures automatic query optimization for the Snowlake Data Cloud architecture.



Query Optimization

- Optimizes query performance
- Analyzes execution plan
- Suggests clustering keys
- Recommends search optimization
- Estimates performance improvement
- Enables continuous improvement
- Provides real-time analytics



Snowpark Python Pipelines


 “Write a Snowpark script for data validation”

 DataFrame transformations

 UDFs for custom logic

 Error handling and logging

 Direct Snowflake connection



Stored Procedures & Tasks

-  “Convert to a stored procedure and schedule it”
-  CREATE PROCEDURE with proper packaging
-  CREATE TASK with WHEN clause
-  Event-driven automation
-  No external orchestrator needed



Subagent Types

Explore — Fast codebase/data exploration

General Purpose — Full tool access

Plan — Architecture and design

Custom — User-defined specialists



Explore — Fast
codebase/data
exploration



General Purpose —
Full tool access



Plan —
Architecture and
design



Custom —
User-defined
specialists



Background Agents

- ❖ Run agents asynchronously
- ❖ Continue working while they execute
- ❖ Monitor with `/agents` or `Ctrl+B`
- ❖ Worktree isolation for safe experiments
- ❖ Kill, resume, or retrieve results anytime



Built-in Skills (40+)

`sql-author` — SQL generation and optimization

`dcm` — Database Change Management

`dbt-projects-on-snowflake` — dbt workflows

`dynamic-tables` — Pipeline creation

`data-quality` — DMF monitoring

`data-governance` — Policies and classification

Custom Skills

-  Markdown files in `.cortex/skills/`
-  Encode team-specific patterns
-  Share via Git or Snowflake stages
-  Combine with `$skill-name` syntax
-  Project, user, or global scope



GitHub Actions Integration



CI: `dbt compile` + `dbt test` on PRs



CD: `snow dbt deploy` on merge to main



OIDC authentication to Snowflake



Dev/prod schema separation



Automated version management



Version Control Flow


1. Code change in Cortex Code
2. `!git commit` with auto-generated message
3. Push triggers CI pipeline
4. Tests pass → merge allowed
5. CD deploys new VERSION\$ to Snowflake



Security & RBAC

- Full Snowflake RBAC enforcement
- OS-level sandboxing
- Three-tier approval system
- Automatic risk assessment
- Never exposes secrets or keys





Docs:

docs.snowflake.com/en/user-guide/cortex-code

dbt on Snowflake:

docs.snowflake.com/en/user-guide/data-engineering/dbt-projects-on-snowflake

Skills:

docs.snowflake.com/en/user-guide/cortex-code/bundled-skills

<https://www.linkedin.com/in/timothyspann/>



Thank You!



<https://github.com/tspannhw/TrafficAI>



THANK YOU



© 2026 Snowflake Inc. All Rights Reserved