



ACCELERATING THE AI AND DATA ENGINEERING LIFECYCLE WITH SNOWFLAKE CORTEX CODE

ENHANCING DEVELOPER EFFICIENCY THROUGH INTELLIGENT CODE GENERATION



Abstract

Artificial intelligence (AI) is redefining development expectations across technology teams in terms of greater speed, higher productivity, and shorter delivery timelines. Tools like GitHub Copilot and Claude have established a new baseline in application development. They support faster code generation, streamline implementation, and shorten development cycles.

AI and data engineering operate in a fundamentally different paradigm. Modern practitioners work across highly fragmented and fast-evolving technology stacks. This breadth significantly increases cognitive load and operational complexity across the entire lifecycle.

Snowflake Cortex Code addresses this challenge with a purpose-built approach. This Whitepaper highlights how Cortex Code delivers context-aware AI support for data engineering workflows. Embedded directly within the Snowflake development environment, Cortex Code understands schemas, pipelines, transformations, and platform-native constructs, providing assistance that is both relevant and actionable, beyond the capabilities of general-purpose tools.

INTRODUCTION

Artificial intelligence (AI) has accelerated the shift toward faster and more efficient software development practices. Teams must now manage data ingestion and transformation through orchestration, quality assurance (QA), and downstream consumption, while also meeting stringent requirements around data governance, security, and privacy.

However, AI and data engineering present a distinct set of challenges. Practitioners work across multiple structured query language (SQL) dialects, Python-based machine learning (ML) frameworks, real-time streaming interfaces, large language models (LLMs), and agent orchestration layers, as well as rigorous governance and compliance controls.

In this environment, generic AI coding assistants often fall short. The most critical challenge is the

context. While these tools excel in standalone coding tasks, they lack awareness of platform-specific data schemas and lineage as well as their underlying governance rules. This results in high-friction debugging across interconnected pipelines and prolonged iteration cycles that delay the delivery of dependable, analysis-ready data.

From a practitioner's perspective, Snowflake Cortex Code accelerates every phase of the AI and data engineering lifecycle. It reduces cognitive overhead, streamlines onboarding, simplifies debugging, and supports faster iteration through intelligent code generation, optimization, explanation, and documentation. As a result, engineering teams can move more efficiently from complex data workflows to trusted, production-ready insights. At the same time, technology leaders can scale delivery velocity across data and AI initiatives.



UNDERSTANDING SNOWFLAKE CORTEX CODE

It is an AI-powered developer assistant that is embedded directly within the Snowflake ecosystem. It acts as a context-aware assistant that understands your data platform, schemas, datasets, and governing role-based access control (RBAC) policies. It accelerates the design, development, and deployment of pipelines across the entire AI and data engineering lifecycle from ingestion, storage, and transformation, to governance, monitoring, and consumption.

Cortex Code provides two distinct developer experiences designed for different user personas:

Snowsight Interface

The Snowsight user interface (UI) is designed for analysts, engineers, and administrators who primarily use the Snowflake platform UI.

The interface requires no setup and provides an interactive experience. It is well-suited for ad hoc SQL development, data exploration, quick prototyping, and account administration through natural language instructions.

Command Line Interface

The command line interface (CLI) supports engineers who prefer working locally in Visual Studio Code (VS Code) or a terminal environment. The CLI brings the capabilities of Cortex Code directly into the local development environment.

With a one-click setup, CLI is best-suited for end-to-end development workflows. It supports Data Build Tool (dbt) projects and pipeline orchestration, while remaining context-aware of

the Snowflake environment. It operates within the boundaries of RBAC permissions, with the ability to connect to external tools through Model Context Protocol (MCP).

Capability Pillars



Build – Accelerated Development and Automation

Focus: Faster creation of data products and pipelines

- Agentic workflow execution enabling autonomous, multi-step task completion
- Context-aware code generation grounded in enterprise schemas and models
- End-to-end pipeline development across ingestion, transformation, and orchestration
- Natural language-to-code (L2C) capabilities that translate business intent into SQL and pipeline logic

Outcome: Significant reduction in development time and onboarding friction



Understand – Contextual Intelligence Across Data

Focus: Greater usability and discoverability of enterprise data

- Contextual data understanding across schemas, lineage, relationships, and semantics
- Intelligent data discovery to identify relevant datasets and possible joins
- Natural language-driven data analysis for rapid exploration

Outcome: Faster insight generation with reduced dependency on undocumented knowledge



Operate – Intelligent Debugging and Operations

Focus: Improved reliability, governance, and operational efficiency

- Intelligent debugging and root cause analysis (RCA) across pipelines
- Snowflake-native operations automation, covering roles, access, and warehouse management
- Workflow observability and issue resolution across interconnected systems

Outcome: Reduced downtime and operational overhead



Optimize – Performance, Cost, and Code Quality

Focus: Continuous improvement of systems and code

- Code refactoring and modernization for legacy SQL and pipelines
- Performance optimization through query tuning and execution efficiency improvements
- Cost optimization through warehouse sizing and resource utilization management
- Best practice enforcement for scalable, maintainable architectures

Outcome: Efficient, future-ready data platforms with controlled costs



CORTEX CODE ACROSS THE AI AND DATA ENGINEERING LIFECYCLE

Snowflake Cortex Code unifies all eight phases of the AI and data engineering lifecycle into a single intelligent, agent-driven system. With the full Snowflake application programming interface (API) coverage, it empowers developers to deliver end-to-end workflows without

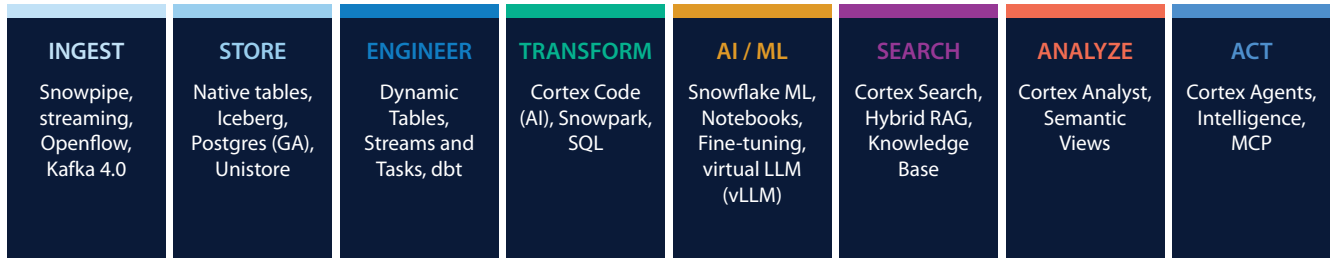
requiring deep specialization across every layer of the stack. This approach creates a multiplier effect, boosting productivity. Table 1 outlines how Cortex Code supports each layer of the AI and data engineering lifecycle, while Figure 1 below illustrates the Infosys accelerators mapped to each phase of the data and AI engineering lifecycle.

Table 1: Cortex Code support across the AI and data engineering stack

| AI and data engineering lifecycle: Cortex Code coverage | |
|---|--|
| Layer One | Ingestion and integration → Snowpipe, streaming, COPY INTO, and schema evolution DDL |
| Layer Two | Transformation and features → Dynamic tables, Snowpark ELT, dbt models, and feature pipelines |
| Layer Three | ML development → Snowpark ML training, Feature Store, Model Registry, and experiment code |
| Layer Four | LLMs and GenAI → Cortex COMPLETE, SUMMARIZE, CLASSIFY, prompt engineering, and prompt chaining |
| Layer Five | Vector Search and RAG → Cortex Search setup, embedding pipelines, and RAG orchestration SQL |
| Layer Six | MLOps and serving → Model deployment UDFs, batch scoring, REST endpoints, and A/B testing |
| Layer Seven | Governance → Masking policies, row-level access, lineage queries, data contracts, and RBAC DDL |
| Layer Eight | Monitoring and optimization → Anomaly detection, cost attribution, quality alerts, and drift detection |

Legend: Snowpark ELT: Snowpark extract, load transform | DDL: Data Definition Language | ML: Machine learning | GenAI: Generative AI | RAG: Retrieval-augmented generation | MLOps: Machine learning operations | UDFs: User-defined functions | REST: Representational state transfer

End-to-end AI and Data Engineering



GOVERNANCE — Horizon Catalog • AI Guardrails • RBAC • Data Masking • Row-Level Security (RLS) • AI Observability • Budgets

Infosys Accelerators Mapped to Each Phase



Fig 1: AI and data engineering lifecycle phases with corresponding Infosys accelerators

Cortex Code in Action

The two examples below demonstrate the practical output that Cortex Code can generate

from a single natural language prompt. Both examples reflect production-oriented workflows rather than simplified demonstrations.

Optimize Mode: Anti-pattern Rewrite

Prompt: Optimize the following query.

ORIGINAL — FLAGGED BY OPTIMIZE MODE

```
SELECT *
FROM raw.events e
JOIN dim.customers c ON UPPER(e.customer_id) = UPPER(c.customer_id)
WHERE YEAR(e.event_ts) = 2026
AND MONTH(e.event_ts) = 4;
```

CORTEX CODE — OPTIMIZED OUTPUT

```
-- Anti-patterns detected:
-- 1. SELECT * on large fact table blocks column pruning
-- 2. UPPER() on join keys prevents partition pruning
-- 3. YEAR()/MONTH() wrapping prevents micro-partition pruning

SELECT e.event_id, e.event_type, e.event_ts, e.event_value,
       c.customer_name, c.segment
FROM   raw.events e
JOIN   dim.customers c ON e.customer_id = c.customer_id
WHERE  e.event_ts >= '2026-04-01'::TIMESTAMP
       AND e.event_ts < '2026-05-01'::TIMESTAMP;
-- Estimated improvement: 60–80% reduction in partitions scanned
```

Generate Mode: RAG Pipeline from Single Prompt

Prompt: Build a Cortex Search RAG pipeline over the support_tickets table.

CORTEX CODE — GENERATED OUTPUT

```
-- 1. Embedding column populated from ticket body
ALTER TABLE support.tickets
  ADD COLUMN body_vec VECTOR(FLOAT, 768);

UPDATE support.tickets
SET   body_vec = SNOWFLAKE.CORTEX.EMBED_TEXT_768('e5-base-v2', body)
WHERE body_vec IS NULL;

-- 2. Cortex Search service over the embedded corpus
CREATE OR REPLACE CORTEX SEARCH SERVICE support.tickets_rag
  ON body ATTRIBUTES product_area, severity, created_at
  WAREHOUSE = SEARCH_WH TARGET_LAG = '1 hour'
  AS SELECT ticket_id, body, product_area, severity, created_at
     FROM support.tickets WHERE created_at >= DATEADD('day', -90, CURRENT_DATE());

-- 3. Retrieval + completion in a single SQL call
SELECT SNOWFLAKE.CORTEX.COMPLETE(
  'mistral-large',
  'Answer using only the context. Q: ' || :question ||
  ' Context: ' || results:context::STRING)
FROM   TABLE(support.tickets_rag! SEARCH(:question, LIMIT => 5)) results;
```

INFOSYS SUCCESS STORIES

- Infosys has combined the strengths of data.ai and Cortex Code to accelerate the enterprise modernization journey of one of the world's largest telecommunications providers. The program extends beyond platform transformation by integrating Infosys accelerators and domain-specific intelligent agents. The focus areas include data modeling, Secure by Design principles, General Data Protection Regulation (GDPR) compliance, and quality engineering. By augmenting Cortex Code with these specialized capabilities, Infosys is enabling faster, more scalable, and governance-driven data product delivery across the enterprise.
- Infosys has used Cortex Code to accelerate the data engineering lifecycle at a global heating, ventilation, and air conditioning (HVAC) company. The integration is helping streamline large-scale modernization initiatives, including Snowpark migration programs. By combining AI-assisted engineering capabilities with Infosys accelerators and best practices, the program enables faster code conversion, improved development efficiency, and a more standardized, scalable approach to modern data platform transformation.
- A leading consumer packaged goods (CPG) company is working with Infosys to leverage Cortex Code and accelerate their end-to-end development lifecycle across modern data and AI initiatives. The partnership integrates Cortex Code with AI-assisted engineering capabilities into the design, development, testing, and deployment processes. The program is driving faster delivery cycles, boosting developer productivity, and promoting greater standardization across enterprise-scale data product development.



BUSINESS IMPACT

Snowflake Cortex Code delivers measurable value across engineering efficiency, infrastructure optimization, and business outcomes. Table 2 summarizes the key business impact dimensions.

Table 2: Business impact dimensions and potential benefits of Cortex Code

| Dimension | Impact Area | Potential Benefit |
|------------------|---|------------------------------------|
| Productivity | Accelerated development and reduced manual effort | 30–50% faster delivery |
| Onboarding | Reduced dependency on undocumented knowledge | 40–60% faster ramp-up |
| Operations | Faster debugging and improved reliability | 25–40% lower issue resolution time |
| Cost efficiency | Query and warehouse optimization | 20–50% compute savings |
| Time to value | Faster data and AI delivery | Accelerated insights |
| Strategic impact | Reduced skill dependency and improved agility | Higher innovation velocity |

These benefits support cost efficiency and faster business value realization, positioning Cortex Code as a strategic enabler rather than just a productivity tool.



ARCHITECTURE AND SECURITY

Every Cortex Code suggestion flows through a multi-stage pipeline. The pipeline includes prompt assembly with live schema context, a fine-tuned Snowflake code-generation model,

syntax validation, schema reference checks, anti-pattern analysis, and a security policy check. The system automatically regenerates suggestions that fail validation. Developers only see code that parses, references real objects, and complies with active masking and row-level access policies.

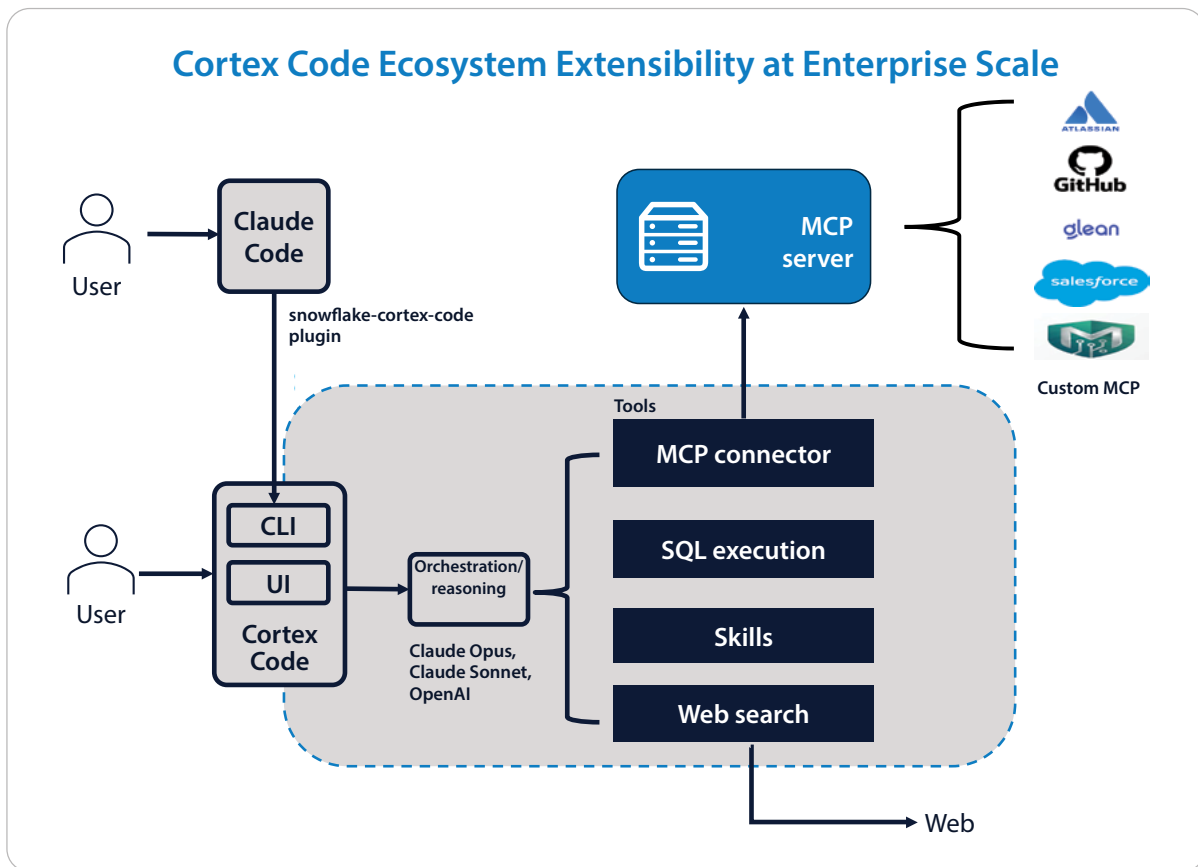


Fig 2: Cortex Code Ecosystem Extensibility at Enterprise Scale



No data egress
Prompts, schema metadata, and generated code remain within the Snowflake cloud environment. The platform does not send any data to external LLM APIs.



Schema metadata only
Cortex Code accesses column names, data types, and constraints. It does not access row-level data.



Role-scoped context
Schema context is filtered based on the user's active role. The model can access only those objects that the user is authorized to query through SELECT permissions.



Audit logging
Every Cortex Code invocation is recorded in ACCESS_HISTORY, providing a complete audit trail for AI-assisted development.

THE FUTURE: TOWARDS AUTONOMOUS DATA ENGINEERING

The Snowflake Cortex Code signals a significant shift toward an autonomous model of data engineering, where intelligence is embedded directly within the platform rather than layered on top.

Looking ahead, data systems will increasingly move toward:



Self-optimizing pipelines that continuously adapt to workload patterns and usage trends



AI-driven orchestration that manages dependencies, execution, and scaling with minimal human intervention



Real-time governance enforcement that applies policies dynamically throughout the development and execution lifecycle

This evolution moves data engineering from manual pipeline construction toward intelligent, self-regulating systems. As a result, engineers can shift their focus from implementation to orchestration and oversight.

CONCLUSION

The Snowflake Cortex Code represents a fundamental evolution in how organizations execute AI and data engineering within the Snowflake ecosystem. By embedding context-aware AI directly into the development environment, the platform helps reduce the gap between intent and production-ready code.

For technology leaders, the implications are clear. Teams can operate with greater efficiency and scale. Development cycles will shrink, governance will be integrated into workflows rather than deferred, and engineers can focus on higher-value architectural initiatives instead of repetitive implementation tasks.

Organizations that adopt Snowflake Cortex Code early through structured enablement, prompt discipline, and governance alignment can strengthen their ability to scale data and AI delivery. This approach will position them to accelerate, innovate, improve operational efficiency, and build greater agility in an increasingly AI-driven landscape.



This Whitepaper has been created and published by Infosys, with inputs from Snowflake.

ABOUT THE AUTHORS



Tom Manfredi

Senior Partner Solution Engineer, Snowflake

Tom Manfredi is a Partner Solution Engineer at Snowflake, where he brings extensive experience in Enterprise Architecture, Data Governance, and Data Management. Tom's career spans four decades in the design and implementation of enterprise-class Data and AI platforms across multiple industries.



Dhanus Kodi Paulraj

Principal Technology Architect, Data Analytics and AI, Infosys

Dhanus Kodi Paulraj is a Principal Technologist and Enterprise Partner Architect for Snowflake at Infosys, where he leads the Snowflake CoE and drives innovation initiatives. He brings deep expertise in architecting large-scale data and AI platforms, enabling enterprises to accelerate transformation and build AI-ready data foundations.



Naveen Kumar Rama

Digital Solution Specialist, Data Analytics and AI, Infosys

Naveen Kumar Rama is a Data Architect at Infosys with deep experience in designing and delivering large-scale data and AI platforms across cloud and enterprise environments. His expertise spans foundational data engineering and next-generation AI, enabling organizations to build robust, AI ready data ecosystems.

Infosys Topaz is an AI-first set of services, solutions and platforms using generative AI technologies. It amplifies the potential of humans, enterprises, and communities to create value. With 12,000+ AI assets, 150+ pre-trained AI models, 10+ AI platforms steered by AI-first specialists and data strategists, and a 'responsible by design' approach, Infosys Topaz helps enterprises accelerate growth, unlock efficiencies at scale and build connected ecosystems. Connect with us at infostopaz@infosys.com.

About Snowflake

Snowflake makes enterprise AI easy, efficient and trusted. Thousands of companies around the globe, including hundreds of the world's largest, use Snowflake's AI Data Cloud to share data, build applications, and power their business with AI. The era of enterprise AI is here. Learn more at snowflake.com (NYSE: SNOW).

For more information, contact askus@infosys.com



© 2026 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/ or any named intellectual property rights holders under this document.