Cost & Performance Optimizer That Turns Snowflake Usage into ROI

A cloud-native solution built on Snowflake Cortex that secures your budget through automated monitoring and actionable agents.

THE CHALLENGE

Unlock greater Snowflake ROI by converting underutilized credits into measurable business value through proactive monitoring of credit trends and usage efficiency.

OUR SOLUTION

GrowthArc's Cost & Performance Optimizer provides an automated, centralized platform designed to help large-scale enterprises gain full transparency over their cloud spend while enabling smarter utilization of credits to drive greater business value.

FinOps Optimization Capabilities

1. Credit Governance & Visibility
2. Warehouse Right-Sizing
3. Storage Cost Recovery
4. Query Performance Optimization
5. Warehouse Queuing Resolution

What Sets It Apart

- Plug-and-play architecture with immediate visibility
- Identifies expensive queries and disk I/O bottlenecks
- Optimizes warehouses and auto-suspend to eliminate waste
- Eliminates congestion through intelligent query routing

Value Delivered

Storage Cost Recovery

Eliminates stale data and high-churn inactive tables for ~35% storage savings

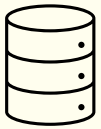

Resource Optimization

Warehouse right-sizing, auto-suspend tuning, 75% dbt refresh reduction, and unused Fivetran pipeline elimination

Comprehensive Spend Tracking

Feature-level insights prevent unmonitored spend escalation

Governance Excellence

80% tagging coverage ensures transparent departmental spend tracking

GrowthArc

Cost & Performance Optimizer: How it Works (Architecture Behind the Scenes)

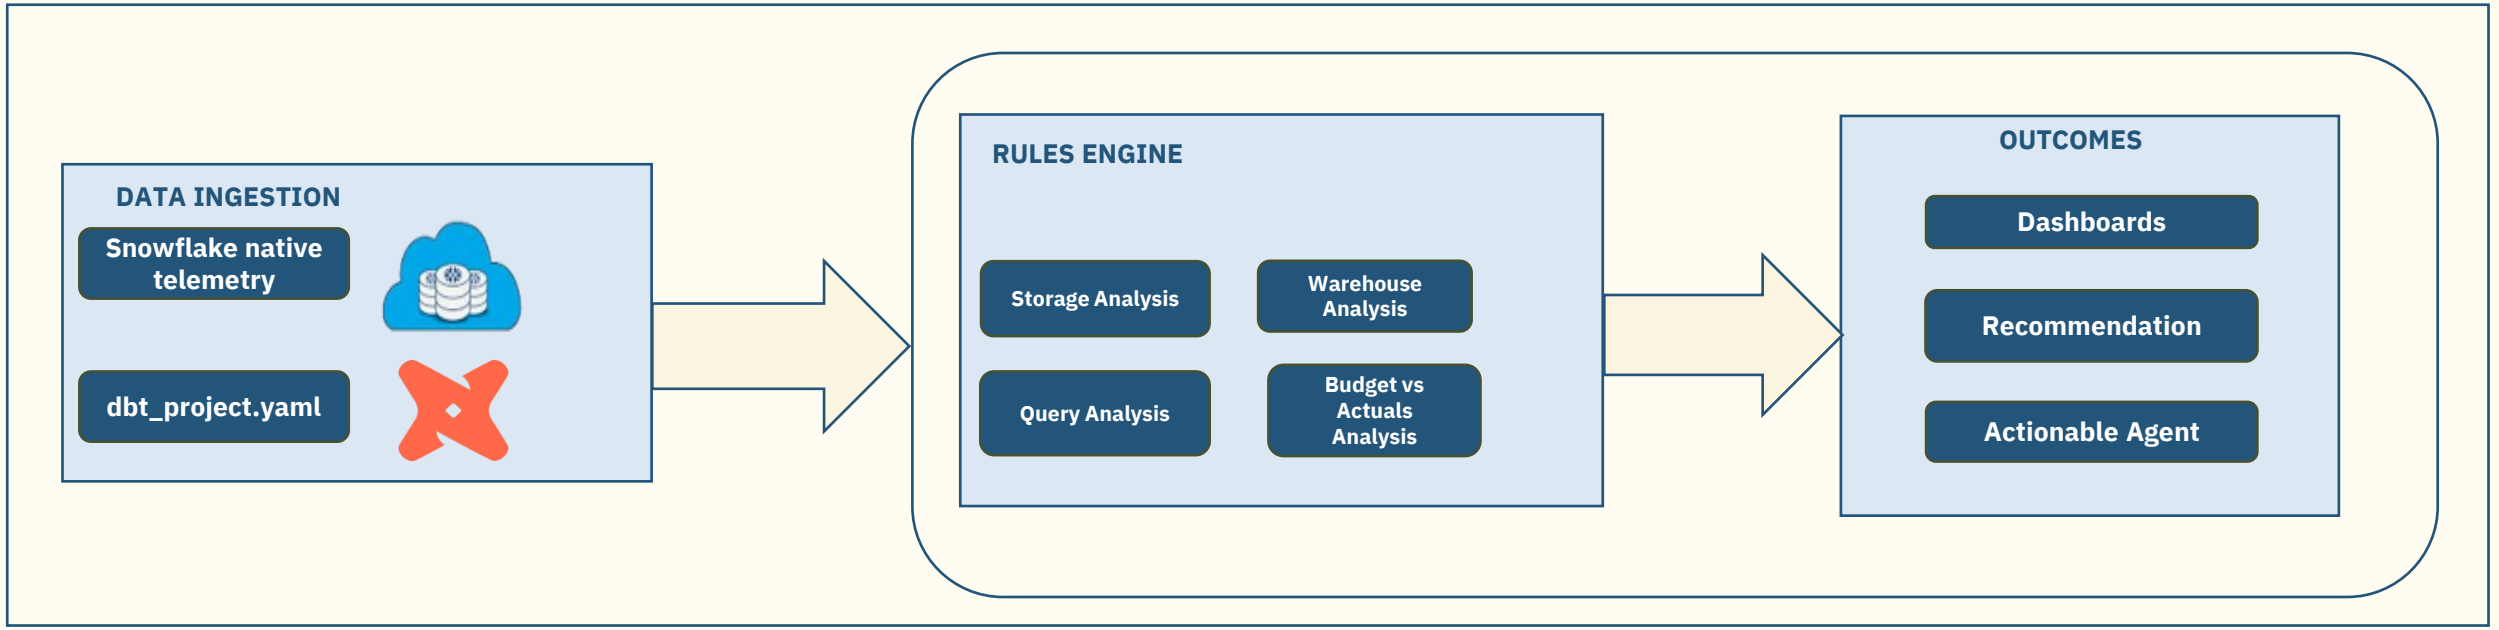

Step 1: Data Ingestion
The foundation of the process involves a plug-and-play architecture with zero data movement.

- **Snowflake Native Telemetry:** The system captures raw performance and usage data directly from Snowflake's native environment.
- **dbt Metadata Integration:** It retrieves dbt models and configuration details from dbt_project.yml to provide context for data transformations.

Step 2: Rules Engine Analysis
The ingested data is processed through a specialized Rules Engine that conducts multi-dimensional analysis:

- **Storage Analysis:** Optimizes Time Travel retention and identifies stale data to recover costs.
- **Warehouse Analysis:** Evaluates compute efficiency by detecting idle time and queuing patterns.
- **Query Analysis**: Pinpoints query inefficiencies and I/O bottlenecks to reduce costs.
- **Budget vs. Actuals Analysis:** Monitors real-time spending against financial guardrails and tagging insights.

Step 3: Intelligence & Governance
The engine identifies specific "Value Leakage" points and provides strategic oversight:

- **Compute Intelligence**: Offers actionable right-sizing recommendations for oversized warehouses.
- **Query Intelligence:** Identifies expensive query joins and excessive disk I/O bottlenecks to optimize spend

Step 4: Strategic Outputs
The final stage transforms analysis into actionable tools for users:

- **Dashboards:** A centralized, self-service platform provides intuitive visualizations of consumption trends.
- **Recommendations:** Delivers specific steps to address identified issues across storage, warehouse, and query categories.
- **Actionable Agent:** Facilitates immediate operational adjustments and proactive governance.

Click Here to Watch Demo **or** Request a Meeting

GrowthArc