

# Snowflake Query Performance Analysis and Optimization

As organizations scale their Snowflake environments, two persistent challenges emerge: unpredictable query performance and spiraling compute costs.

Long-running queries, over-provisioned warehouses, idle clusters, and inefficient SQL patterns can quietly [drain millions of Snowflake Credits each year](#). This is precisely where Revefi's AI Agent steps in. Not just as a monitoring tool, but as an auto-pilot or co-pilot that actively manages and optimizes Snowflake query performance, and warehouse utilization in real time.

## Solving Long-Running Queries: Optimization for Peak Performance

One of the most common reasons behind Snowflake inefficiencies is long-running queries. This often stems from performance bottlenecks like excessive data scans, inefficient joins, or redundant operations, wasting Snowflake credits and stalling downstream analytics pipelines.

Revefi's AI Agent (RADEN) provides real-time detection and automated remediation, [making Snowflake query optimization seamless and instantaneous](#). RADEN solves this by:

- Scanning every query execution to pinpoint issues such as high queuing times or suboptimal filtering.
- Generating optimized query rewrites tailored to Snowflake's architecture.

The screenshot displays the Revefi AI Agent interface. On the left is a navigation sidebar with options like Feed, Exec Summary, Top Questions, Monitors, Data, Views, Warehouses, BI Artifacts, Pipelines, Usage, Workflows, Tags, Connections, and Tickets. The main panel shows a 'Query Run' summary for Query ID: 01bf67af-0206-8f88-0006-8c0708b16586. Key metrics include: Query Type: SELECT, User Name: DEMO\_USER, Warehouse Name: REVEFI\_DEV\_WH, Query Parameterized Hash: 51a842..., Warehouse Size: xs, Query End Time: 30 Sep 2025 5:47 PM, Bytes Scanned: 35.88 GB, Rows Produced: 33, Percentage Scanned From Cache: 31.8%, Total Elapsed Time (s): 11 min 29 sec, Execution Time (s): 1 min 48 sec, Compilation Time (s): 3 sec, Queued Provisioning Time (s): 0 sec, Queued Overload Time (s): 9 min 40 sec, Execution Status: Success.

The 'Query Text' section shows a complex SQL query with multiple joins and filters. An AI-generated 'Query Analysis' window is overlaid on the right, providing detailed insights:

Based on my analysis of the query and its execution plan, here's an optimized version with detailed explanations:

### Query Analysis

The original query (01bf67af-0206-8f88-0006-8c0708b16586) took **11.5 minutes** to execute, with significant queuing time (9.7 minutes). It scanned **38.5 GB** of data and processed 5,398 rows after the join.

### Key Performance Issues Identified:

1. **High queuing time** (579 seconds) - warehouse REVEFI\_DEV\_WH was overloaded
2. **Large data scan** (38.5 GB) with low cache hit rate (31.8%)
3. **Redundant filtering** on both tables for the same conditions
4. **Inefficient join** between large account usage tables

### Optimized Query

```
WITH filtered_query_history AS (
  SELECT
    query_id,
    query_parameterized_hash,
    warehouse_name,
    query_type,
    query_text,
    execution_time
  FROM
    "SNOWFLAKE".account_usage.query_history qh
  WHERE
    qh.warehouse_name IS NOT NULL
    AND qh.warehouse_name IS NOT NULL
    AND qh.query_parameterized_hash IS NOT NULL
    AND qh.query_parameterized_hash IS NOT NULL
    AND qh.execution_time > 0
    AND qh.start_time >= '2025-09-23 00:00:00.0' -- up to start_time, inclusive
    AND qh.start_time < '2025-09-30 21:35:48.700878297' -- up to the end_time, non-inclusive
)
```

At the bottom, there is an 'Add Context' button and a chat input field with a 'Claude Sonnet 4' model selector.

For instance, it applies advanced techniques like cardinality reduction (pruning unnecessary data paths early), and eliminates redundant filters that inflate scan volumes.

### Results:

The benefits are immediate and measurable. Data teams using RADEN report significant reductions in compute waste, with query execution times dropping dramatically, which translates to lower overall Snowflake costs (often by 40-60% for query-heavy workloads).

# Handling Workload Spikes: Intelligent Scaling to Prevent Cost Overruns

Workload spikes add another layer of complexity to Snowflake query performance management. These bursts (often triggered by concurrent BI reports, ad-hoc analytics, or seasonal demands) can overwhelm warehouses, causing queues to balloon and queries to crawl.

Without intervention, they drive up Snowflake credit consumption unpredictably, eroding ROI on your data platform.

RADEN's strength in this area lies in its predictive analytics and user-centric insights. The AI Agent:

- Dives deep into Snowflake's usage patterns, identifying the exact users or roles responsible for spikes.
- Correlates high credit burn with specific times of day or activities, such as a marketing team dashboard refresh at 9 AM or an engineering ETL job at midnight.

The screenshot shows the Raden AI Agent interface. The main window displays a report titled "Here's what I found regarding spikes in Snowflake spend tied to users, workloads, and times of day:". The report is structured as follows:

### 1. Spikes by User and Workload (Warehouse)

- The largest credit consumption in the last 7 days is concentrated among a few users and warehouses:
  - User DEMO\_USER:
    - PC\_DBT\_WH: 84.87 credits
    - REVEFI\_DEV\_WH: 62.85 credits
    - REVEFI\_WH: 0.09 credits
  - User REVEFI\_USER\_BETA:
    - REVEFI\_WH\_BETA: 49.95 credits
  - User PC\_DBT\_USER:
    - PC\_DBT\_WH: 9.29 credits
    - PC\_DBT\_DEMO\_WH: 4.79 credits
    - PC\_DBT\_STAGING\_WH: 4.60 credits
    - PC\_DBT\_DEV\_WH: 4.62 credits
    - PC\_DBT\_TEST\_WH: 4.61 credits
  - User SYSTEM:
    - LOOKER\_WH: 15.74 credits
    - TABLEAU\_WH: 14.02 credits
    - SLA\_MONITOR\_WH: 10.28 credits
    - REVEFI\_WH\_BETA: 5.03 credits

### 3. Summary Table: Top Spenders (Last 7 Days)

User	Warehouse	Credits (7d)	Notable Query Spikes (per day)
DEMO_USER	<a href="#">PC_DBT_WH</a>	84.87	20,857 (2025-09-24)
DEMO_USER	<a href="#">REVEFI_DEV_WH</a>	62.85	26,100 (2025-09-24)
REVEFI_USER_BETA	<a href="#">REVEFI_WH_BETA</a>	49.95	16,167 (2025-09-24)
SYSTEM	<a href="#">LOOKER_WH</a>	15.74	90 (multiple days)
SYSTEM	<a href="#">TABLEAU_WH</a>	14.02	105 (multiple days)
SYSTEM	<a href="#">SLA_MONITOR_WH</a>	10.28	406 (multiple days)
PC_DBT_USER	<a href="#">PC_DBT_WH</a>	9.29	1,037 (2025-09-24)

### 4. Key Takeaways

- Spikes in spend are driven by a handful of users and their workloads.
- Spikes are highly temporal — certain days see massive query volumes and credit usage, while others are much lower.
- Warehouses with the highest spend are [PC\\_DBT\\_WH](#), [REVEFI\\_DEV\\_WH](#), and [REVEFI\\_WH\\_BETA](#).

Data Sources

- Credit usage: AccessedQuervCountByUserMaterialized (last 7 days)

From this, RADEN generates dynamic reports, including ranked lists of top spenders over customizable periods, empowering admins to enforce governance without guesswork.

What sets RADEN apart is its agentic nature. It reshapes workflows instantly by recommending proactive adjustments, like auto-scaling warehouses during predicted peaks or isolating spiky users to dedicated clusters.

## Results:

In [Revefi's workload spike video](#), the Revefi Agent AI uncovers patterns like "a single user driving 40% of evening consumption," enabling targeted interventions that curb spikes before they escalate.

This not only stabilizes query performance (reducing average latency during peaks), but also slashes costs by optimizing resource allocation in real time.

# Use The Skills of an AI Agent That Continuously Learns

By integrating long-running query optimization with workload spike management, RADEN delivers a holistic approach to Snowflake performance. It learns from historical and real-time data, compounding intelligence over time to anticipate issues. For example, if spikes correlate with unoptimized queries, RADEN prioritizes rewrites for those patterns, creating a feedback loop of continuous improvement.

Data teams gain a unified dashboard for visibility (query heatmaps, credit trends, and optimization recommendations), all actionable via natural language prompts.

In an era where data velocity defines success, Revefi's AI Agent stands out as the ultimate Snowflake expert.

## Ready to Optimize Your Snowflake Query Performance and Costs?

Revefi's AI Agent delivers the only fully autonomous solution that combines deep visibility, predictive scaling, and continuous query optimization into a single platform. Stop overpaying for Snowflake and start running the leanest, fastest Snowflake environment possible.

[Talk to an expert](#), or schedule a demo today to see your personalized savings projection in under 15 minutes.

COOL VENDORS is a trademark of Gartner Inc. and its affiliates. Gartner does not endorse any company, vendor, product, or service depicted in its publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner publications consist of the opinions of Gartner's business and technology insights organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.

**Gartner**

**COOL  
VENDOR  
2025**

## Why Leading Companies Choose Revefi for Snowflake Optimization



**SOC 2  
TYPE II  
CERTIFIED**

---

## Integrations



**databricks**



Google  
Big Query



**amazon**  
REDSHIFT

---

## Organizations That Trust Revefi



*StanleyBlack&Decker*



*Ingersoll Rand*



**Verisk**



**Cribl**

