REVEFI

Solution Brief

# AI Agent for Google BigQuery

Gartner COOL VENDOR 2025

## Revefi's Unified FinOps, DataOps, and Data Observability

By combining cost control, operational resilience, and data trust, Revefi transforms BigQuery into a strategic asset rather than an unpredictable expense line item.

- Eliminate wasted BigQuery spend, identify inefficient queries and anomalies, and deliver AI-driven optimization recommendations that cut cloud data costs
- Detect schema drift, stale tables, and performance issues early to get instant root-cause analysis (RCA) and quality alerts
- Correlate cost, performance, and data health in one view to enforce proactive, automated governance without slowing innovation

## How Can Organizations Leverage Revefi?

Revefi retrieves metadata from Google BigQuery in minutes, requiring no code changes or complex integrations. It continuously monitors query patterns, slot usage, table access, and cost drivers at a granular level, delivering real-time visibility across projects, datasets, and teams.

**1. Automated Cost Management, and Optimization**

Revefi delivers comprehensive slot and cost visibility across BigQuery workloads by analyzing query activity, bytes processed, and reservation slot usage. The platform identifies expensive query patterns, inefficient pipelines, and unused resources while providing optimization insights.

**Key capabilities:**

- Unified visibility into BigQuery slot consumption patterns and cost drivers
- Identification of high-cost queries by slot-ms usage, bytes processed, and runtime
- Automatic comparison of utilization across reservations (dev, prod, QA)

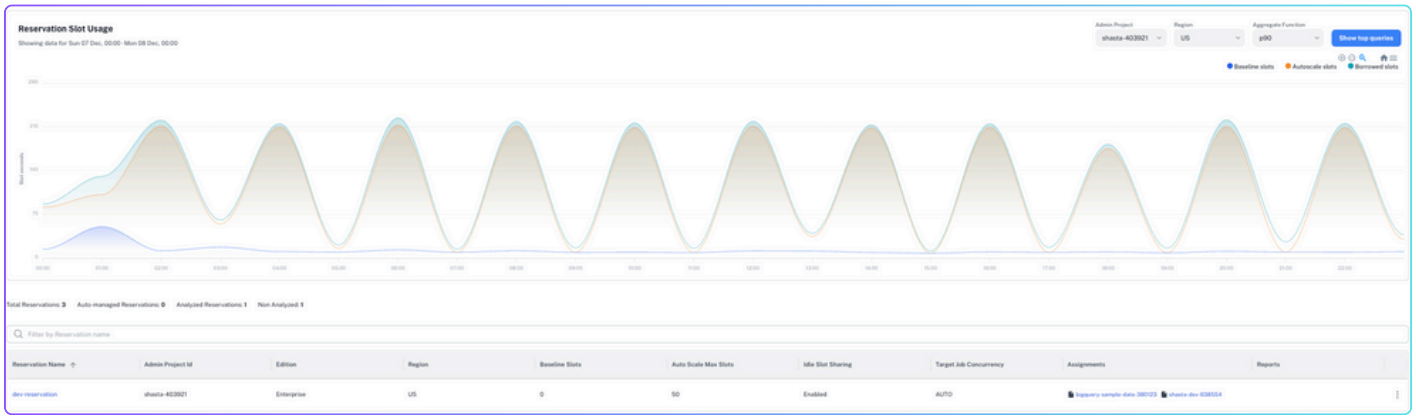

### 2. Automated Performance and Efficiency Insights

Using metadata from and job execution logs, Revefi surfaces trends in execution time, queue time, job concurrency, throughput and more. This allows data teams to diagnose bottlenecks such as:

- Overloaded reservations
- Queries with unusually high slot consumption
- Long-running operations

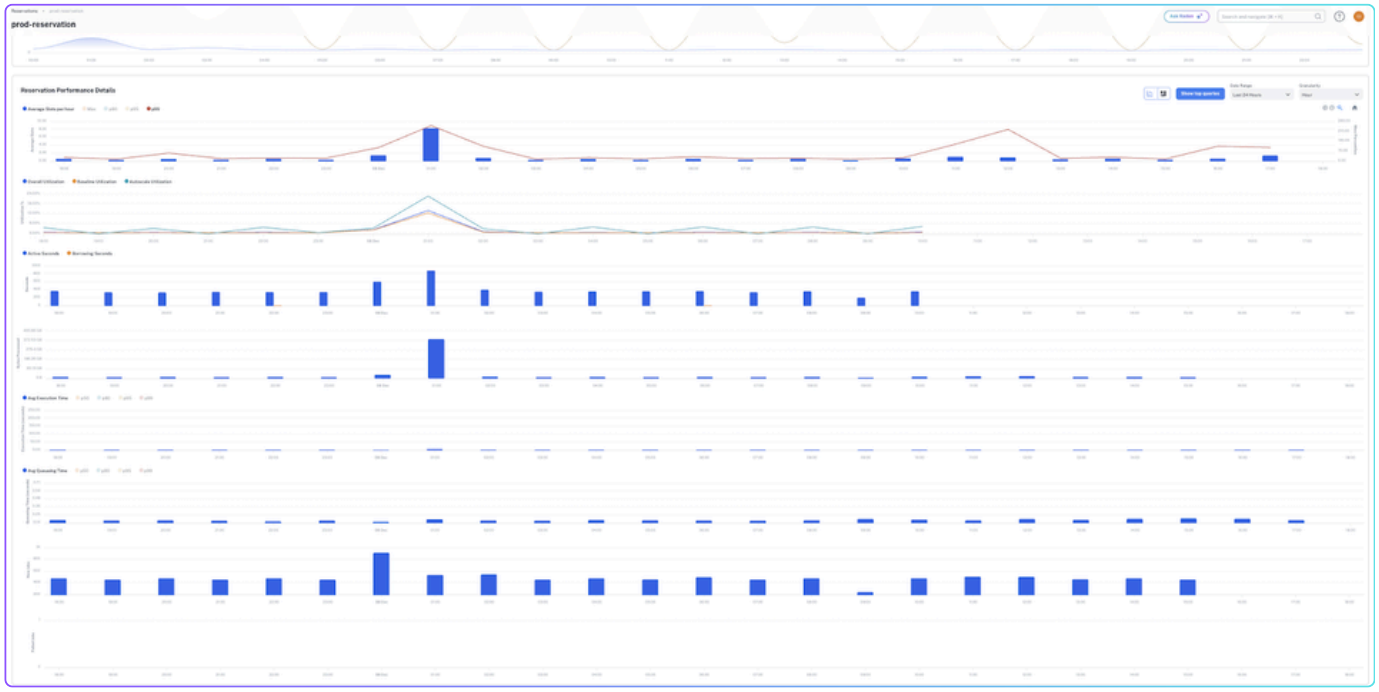


### 3. Automated Data Quality, and Observability

Revefi automatically generates and maintains data quality monitors across every table in an organization's BigQuery environment. These monitors track schema drift, data freshness, load success/failure, row count anomalies, and threshold-based custom SQL rules.

**Our platform provides:**

- Daily checks (see image below)
- Detailed trend lines for each monitor
- Custom SQL monitors to enable high-sensitivity anomaly detection

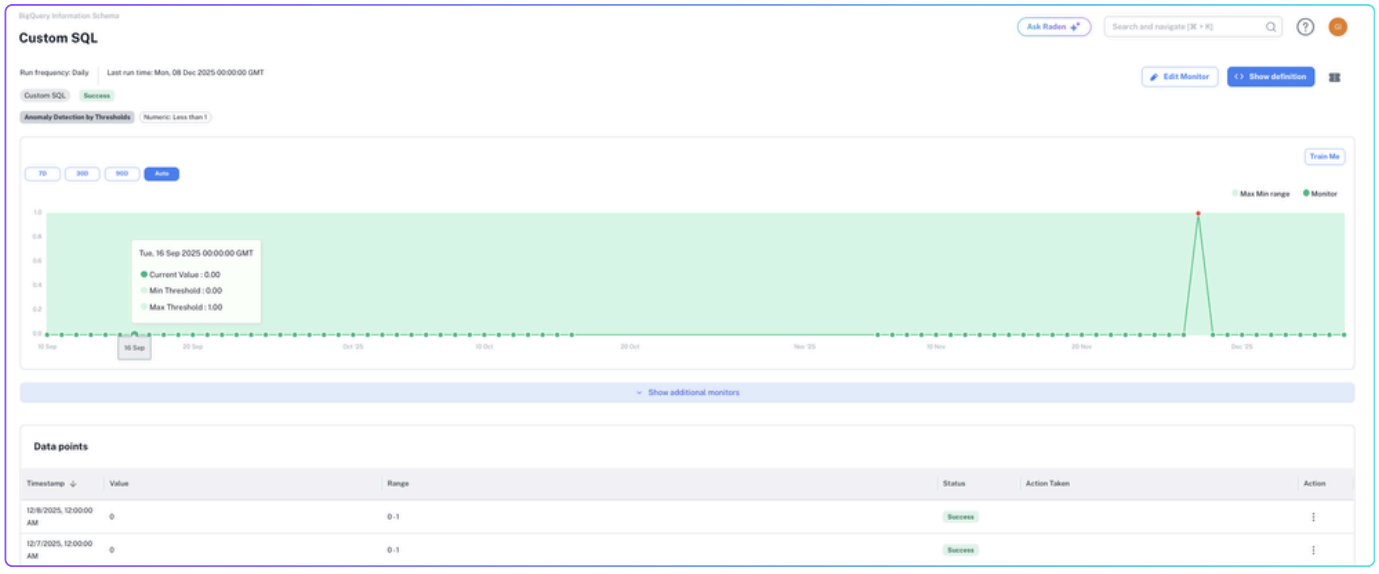

### 4. Granular Data Usage, and Access Intelligence

Revefi provides granular insights into BigQuery compute costs, broken down to individual users, tables, and workloads. This makes it possible to understand how data assets are being accessed, which workloads are responsible for cost spikes, and which teams are under or over-utilizing compute resources.

User-level, and job-level breakdowns include:

- Slots consumed
- Bytes processed
- Runtime and queue time
- Behavior analysis

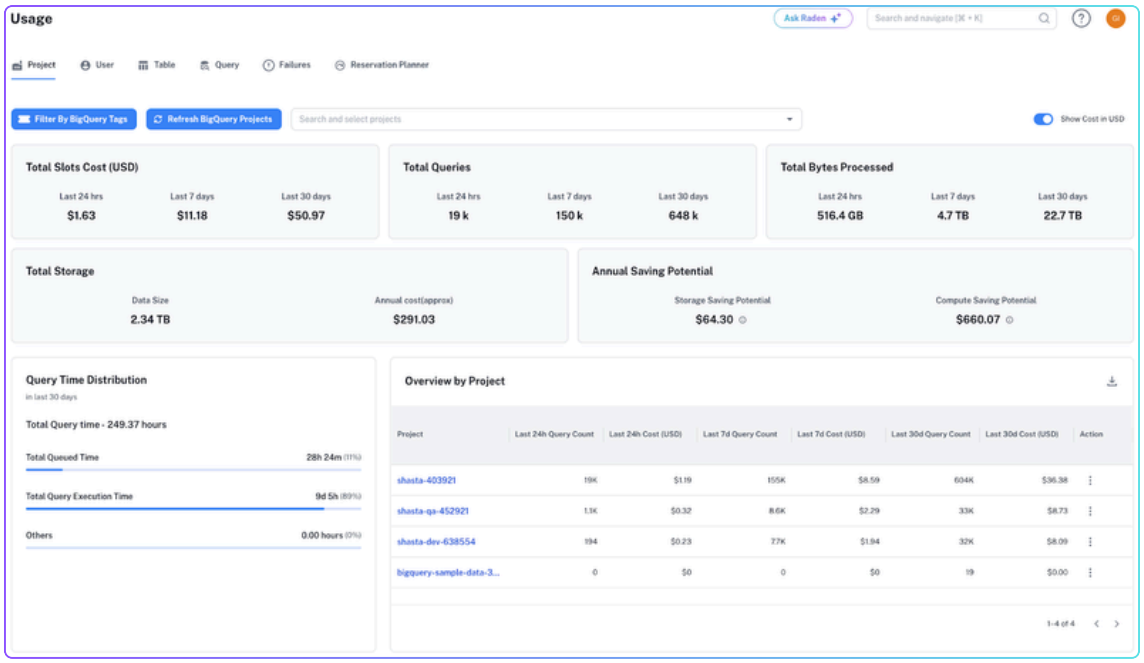


### 5. Pipeline and Orchestration Observability

Revefi integrates with orchestration systems such as Cloud Composer (Airflow). It monitors DAG state, execution duration, queueing patterns, and success/failure trends. Engineers gain a unified view of operational pipeline health across BigQuery and workflow systems.

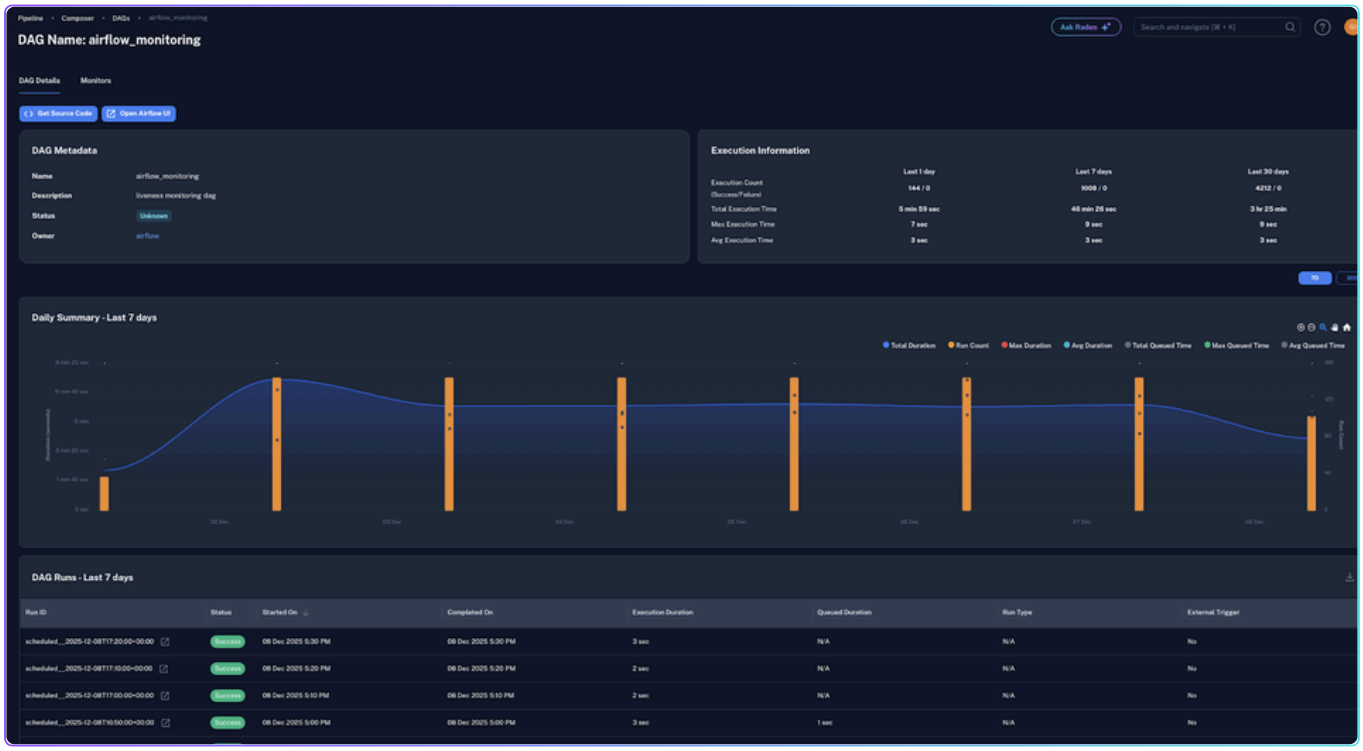

### 6. BigQuery Table Health, Lineage, and Metadata Intelligence

Table-level dashboards show daily incremental row counts, load durations, total load bytes, and schema metadata. Revefi's lineage visualization reconstructs upstream and downstream dependencies using BigQuery job history, enabling impact analysis during failures or schema changes.

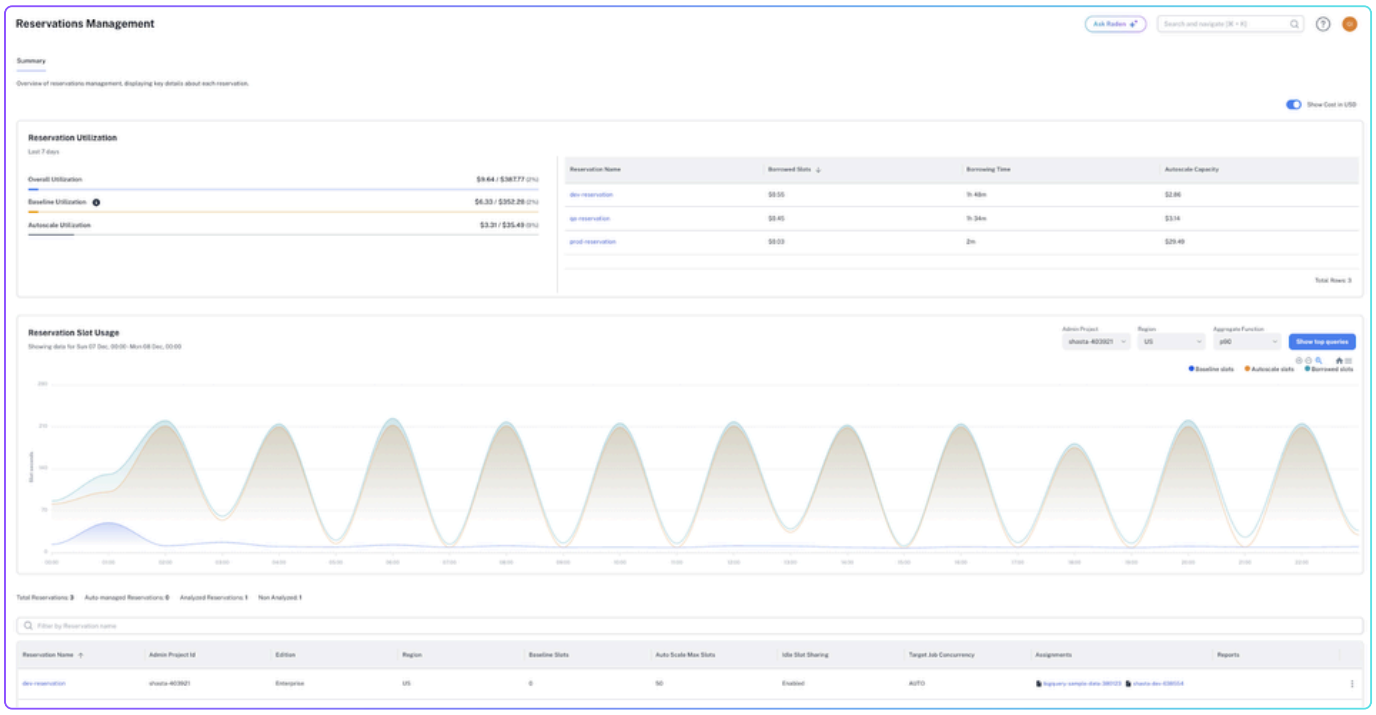


## Put Google BigQuery on Autopilot with Revefi's Autonomous AI

With Revefi, data teams now eliminate cost surprises, are able to debug slow dashboards, reverse-engineer why data pipelines silently broke, continuously for every query, slot, table, and job, and schema change across all projects.

## Integrations

snowflake databricks Google Big Query amazon REDSHIFT

## Organizations That Trust Revefi

AMD StanleyBlack&Decker Ingersoll Rand Ocean Spray Verisk Cribl The Princeton Review