

# One of the largest Google BigQuery Enterprises in the US runs on Revefi

The data team made Revefi the operating system for their BigQuery environment for spend, capacity planning, observability, and DataOps.

## Revefi platform is now the default to manage BigQuery.

The customer, one of the largest BigQuery footprints in the US, was scaling fast with more projects, more Looker users, and more pipelines and were unable to keep up. Spend was a leaky black box with no way to trace BigQuery costs back to Looker or business users driving them. A major cloud migration and a contract renewal were forcing a high-stakes capacity decision.

The data platform team was the last to know when something broke, failures arrived as stakeholder Slack messages, not alerts. Instead of building more internal tools, the team consolidated on Revefi across all of these areas at once, deploying iteratively and shaping each capability from slot-optimization logic to alert routing around their own requirements.

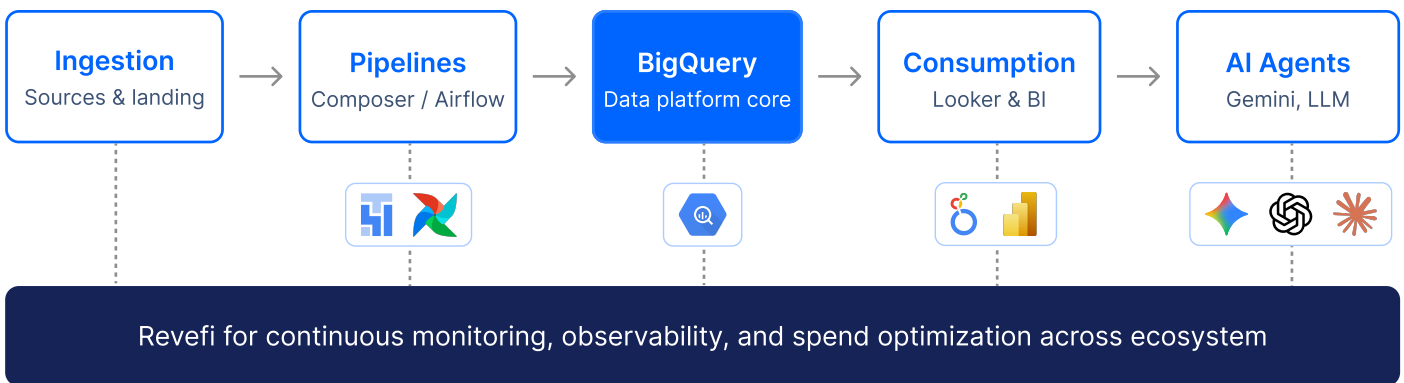


Figure 1: Revefi manages the entire BigQuery data stack, continuously monitoring spend, performance, SLAs, and lineage

<p><b>&gt;70x</b></p> <p>Growth in table SLA monitors in under two quarters</p>	<p><b>~6x</b></p> <p>Increase in total active monitors</p>	<p><b>~3.5x</b></p> <p>more tables under continuous monitoring</p>
<p><b>~2x</b></p> <p>Growth in monitored GCP projects</p>	<p><b>3x</b></p> <p>Expansion in connected systems: BigQuery, Looker, Composer &amp; Gemini</p>	<p><b>100%</b></p> <p>Data-SLA alerting &amp; RCA routed through one platform</p>

## Facing the Obstacles: What Were the Biggest Challenges?

Historically, the organization has faced significant challenges when it came to managing cost, and data accountability across their Google BigQuery ecosystem.

### **Spend was a leaky black box:**

No way to trace BigQuery costs back to Looker or the business users running them.

### **Lack of details on Capacity Usage:**

During migration and contract renewal, it lacked fine-grained usage data to baseline decisions.

### **Multiple monitoring tools:**

The three monitoring tools were inadequate with gaps, none connected.

### **Unable to handle incidents from the business:**

Pipeline failures arrived as stakeholder Slack messages without any automated detection and no alert routing.

## With Revefi: Four Use Cases, One Platform.

### 01 BigQuery Spend Metering: Costs are traceable and optimized

Revefi built end-to-end lineage from Looker explores to BigQuery queries, so every expensive query traces to a specific explore and a specific user and business teams see the cost of what they run.

- Redundant Looker dashboard queries identified and eliminated with their business owners.
- When a new BI tool was onboarded, its slot (and dollar) spike was caught and resolved within the billing cycle not after.
- Every future tool or workload onboarded to BigQuery inherits the same coverage automatically.

### 02 Capacity Planning: Fully defensible commitment

With a regional migration and a Google contract renewal underway, the team faced one question: how many slots to commit versus autoscale? Revefi became the system of record for that decision and the place the team now answers “what if we change the committed slot count?” in seconds.

- Cross-reservation slot analysis down to second-level granularity, across reservations and projects.
- Historic usage, slot oversubscription, and predicted demand modeled to strike the right balance of performance and spend.
- Phased to match the migration timeline, with what-if analysis on committed-slot changes driven entirely from Revefi.
- The recommended commit-vs-autoscale balance went straight into the Google contract negotiation, grounded in data, not estimates.

### 03 DataOps: Pipeline observability and alert routing

Pipeline alerting lived in one legacy tool while custom scripts tracked costs none of it integrated. Revefi replaced that patchwork and is now the single destination for pipeline monitoring, root-cause analysis, and alert routing across the organization.

- Composer and Airflow pipelines monitored for spend, performance, and SLA status.
- Alerts route to the team that owns the data via Slack with the data platform team no longer the manual middleman.
- Full lineage from pipeline to BigQuery tables to Looker explores.
- Centralized alert-management integration underway as a joint build.

### 04 Data Observability: Automatically catching what rules miss

Legacy tooling couldn't catch a pipeline that ran without error but produced no data. Revefi auto-configures freshness thresholds per table from historical behaviour, and SLA-monitoring adoption grew rapidly as teams saw it work.

- Looker usage reporting migrated out of Excel: historical trends, anomaly detection, peak-time analysis.
- Revefi sets freshness thresholds per table automatically, with no manual setup.
- Query-efficiency and utilization metrics surface underperforming workloads before they hit the business.
- SLA monitoring expanded from a small manual deployment to broad coverage of critical tables.

## What's next?

The team is extending spend visibility and operational control built for BigQuery into a broader data and AI estate, including Databricks and AI FinOps. As AI-driven pipelines and LLM-based agents take on more of the consumption layer, the same questions will surface, and Revefi can already answer them.

## 01 BigQuery Spend Metering: Costs are traceable and optimized

Revefi built end-to-end lineage from Looker explores to BigQuery queries, so every expensive query traces to a specific explore and a specific user and business teams see the cost of what they run.

- Redundant Looker dashboard queries identified and eliminated with their business owners.
- When a new BI tool was onboarded, its slot (and dollar) spike was caught and resolved within the billing cycle not after.
- Every future tool or workload onboarded to BigQuery inherits the same coverage automatically.

## 02 Capacity Planning: Fully defensible commitment

With a regional migration and a Google contract renewal underway, the team faced one question: how many slots to commit versus autoscale? Revefi became the system of record for that decision and the place the team now answers “what if we change the committed slot count?” in seconds.

- Cross-reservation slot analysis down to second-level granularity, across reservations and projects.
- Historic usage, slot oversubscription, and predicted demand modeled to strike the right balance of performance and spend.
- Phased to match the migration timeline, with what-if analysis on committed-slot changes driven entirely from Revefi.
- The recommended commit-vs-autoscale balance went straight into the Google contract negotiation, grounded in data, not estimates.

## 03 DataOps: Pipeline observability and alert routing

Pipeline alerting lived in one legacy tool while custom scripts tracked costs none of it integrated. Revefi replaced that patchwork and is now the single destination for pipeline monitoring, root-cause analysis, and alert routing across the organization.

- Composer and Airflow pipelines monitored for spend, performance, and SLA status.
- Alerts route to the team that owns the data via Slack with the data platform team no longer the manual middleman.
- Full lineage from pipeline to BigQuery tables to Looker explores.
- Centralized alert-management integration underway as a joint build.

## 04 Data Observability: Automatically catching what rules miss

Legacy tooling couldn't catch a pipeline that ran without error but produced no data. Revefi auto-configures freshness thresholds per table from historical behaviour, and SLA-monitoring adoption grew rapidly as teams saw it work.

- Looker usage reporting migrated out of Excel: historical trends, anomaly detection, peak-time analysis.
- Revefi sets freshness thresholds per table automatically, with no manual setup.
- Query-efficiency and utilization metrics surface underperforming workloads before they hit the business.
- SLA monitoring expanded from a small manual deployment to broad coverage of critical tables.