yuki

CUSTOMER CASE STUDY

Riskified Cuts Looker Compute Spend Approximately 50% With Yuki

Real-time query routing for self-service analytics on Snowflake



Riskified | Snowflake | Looker

About Riskified

Riskified empowers ecommerce growth by outsmarting risk. Fraud, policy abuse, and false declines don’t just threaten security. They also cut into revenue, stretch teams, and erode margins. As fraud grows faster and more sophisticated, merchants need more than a tool. They need a partner accountable for the outcome.

That's the difference Riskified brings. Its AI-powered fraud and risk intelligence platform draws on machine learning trained across a global network which based on data available to us, includes approximately 50 publicly held companies worldwide. This gives merchants the accuracy to approve more good orders, block fraud, and curb abuse without adding friction for legitimate customers.

Riskified also stands behind its decisions. Merchants can shift chargeback liability and calibrate performance to match their own risk tolerance and business goals. The result is less uncertainty and a clearer path to profitable growth.

For merchants under pressure to grow revenue while managing risk, Riskified turns fraud prevention from a cost center into a driver of confident, sustainable growth.

DATA STACK

- Snowflake
- Looker
- dbt
- Airflow
- Databricks
- Apache Iceberg



Challenges

01 One warehouse size could not fit every query

Most queries in Looker required relatively little compute, while a smaller share needed significantly more. Sizing the warehouse for the largest queries led to overspending on most workloads. Sizing it for the average query risked slowing down the requests analysts cared about most.

The default warehouse was built oversized by design.

02 Ad hoc workloads were impossible to predict

Riskified could profile scheduled ETL jobs and match them to the right warehouse sizes.

Analyst workloads were different. The data team had no reliable way to know which query would run next, how much data it would scan, or how much compute it would need, making manual resizing impractical at that pace.

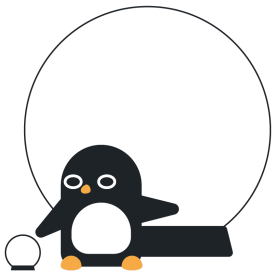
03 Analysts should not have to think about infrastructure

Riskified's analysts focus on fraud prevention and risk signals, as well as related, time-sensitive business questions. Warehouse sizing is not their job.

The team needed a way to control Snowflake costs without asking analysts to choose compute resources, change connection settings, or understand the infrastructure behind every query.

Solution

Riskified deployed Yuki to make warehouse-sizing decisions automatically for Looker and for selected Snowflake workloads.



01 A real-time routing layer

Yuki sits in the query path and evaluates each workload as it arrives.

Instead of sending every Looker query to the same default warehouse, Yuki routes the query to the warehouse size that matches its actual compute requirements.

02 Right-size every query, not the average query

Lightweight queries can run on smaller, more efficient warehouses.

Heavier queries still receive the compute they need.

03 Invisible optimization

Analysts continue using Looker as usual. They do not need to choose warehouse sizes or think about Snowflake infrastructure.

The data team does not need to continuously review queries or manually resize warehouses.

Results

~50%

Looker compute cost reduction

Riskified achieved approximately 50% in savings across Looker workloads managed by Yuki.

Looker was the hardest workload to optimize manually, and that's exactly where automated routing delivers the most value.

Most queries did not need the warehouse that they would have used by default. Yuki identified the lighter ones and routed them to more efficient compute automatically. The savings were consistent from day one.

~20%

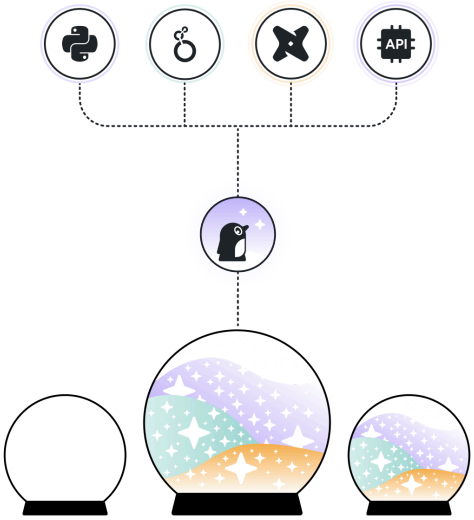
Broader workload savings

Across the broader mix of workloads routed through Yuki, including more predictable ETL traffic, Riskified observed approximately 20% savings.

Predictable ETL can often be profiled and manually optimized.

Ad hoc analytics changes from one query to the next.

The more unpredictable the workload, the more value Riskified received from automated routing.



Results

Hands-off optimization

For Ran Salmona, VP Data at Riskified, not needing to check the dashboard was a sign that the system was working.

No analyst training required

Riskified's analysts continued to ask business questions while Yuki automatically handled the infrastructure decision.

There was no need to train users on warehouse sizing or ask them to decide how much compute a query might need.

A clear qualification question

Riskified now uses a simple question when discussing Yuki with other Snowflake customers:

How much of your Snowflake traffic comes from analysts?

When most traffic comes from scheduled ETL, manual tuning may solve part of the problem.

When a large percentage comes from analysts, Looker, and self-service BI, static warehouse sizing becomes difficult to manage efficiently.

That is where Yuki delivers the most value.

Ad hoc analytics is where Yuki gives us the most value. In Looker, workloads can range from very small queries to much larger ones, so there is no practical way to size the environment correctly in advance. Yuki solves a problem I do not think we could solve manually.



Ran Salmona
VP Data @ Riskified

> “If most of your traffic comes from analysts, you need a solution like Yuki. There is no practical way to manually optimize a workload that changes from one query to the next.”
>
> Ran Salmona, VP Data @ Riskified

Riskified's Looker environment shows why static warehouse sizing breaks down for self-service analytics. Most queries do not need large compute, but the few that do force teams to overprovision everything. Yuki makes that decision query by query, cutting waste without putting infrastructure decisions on analysts.



Ido Arieli Noga
CEO and Co-Founder @ Yuki

yuki



Get a free analysis!

Find out how much you can save

Free Analysis

yukidata.com/free-analysis