

stripe

Inspecting Trino on Ice

Kevin Liu | Stripe

@kevinliu



What is Stripe?

stripe



Payments infrastructure for the internet

Millions of companies of all sizes—from startups to Fortune 500s—use Stripe's software and APIs to accept payments, send payouts, and manage their businesses online.

Start now >

Contact sales >



Atlas



Capital



Climate



Corporate
Card



Identity



Issuing



Payment
Links



Radar



Sigma



Terminal



Data
Pipeline



Billing



Checkout



Connect



Elements



Invoicing



Payments



Payouts



Revenue
Recognition



Tax



Treasury



Financial
Connections

stripe

What is Trino?

- ✔ Highly parallel and distributed query engine
- ✔ Enables interactive data analytics
- ✔ Data Lakehouse architecture



Trino @ Stripe

Internal use cases

- ✔ Ad hoc analytics
- ✔ Dashboards
- ✔ API for internal services and data apps
- ✔ No ETL (yet)

Internal use cases

Trino is the foundational infrastructure on which other apps and services are built upon.

“swiss army knife” in data ecosystem - me, 2023

Internal use cases

Services

- **Hubble (internal, similar to Superset): ad hoc analytics**
- **Notebooks: data analysis**
- **Streamlit: data visualization**
- **Checker (internal): define and run integrity checks against data**

Internal use cases

Data Apps

- **Inspect Iceberg Tables ([go/iceberg/info](#))**
- **Inspect Table Metadata from Hive Metastore**
- **Inspecting Iceberg usage with Trino queries**

Inspecting Iceberg Tables ([go/iceberg/info](https://go.iceberg/info))

Problem with Iceberg CLI

- Requires mydata
- Requires bazel build (~30 seconds)
- Single command only (`show-metadata / show-partitions`)

😬 `pay exec bazel run`
`//src/scala/com/stripe/iceberg/cli:cli` -

- 🤖 Machine friendly, not human friendly

Trino + Iceberg!

Trino Iceberg connector and metadata tables!

Same information!

`$history`

`$snapshots`

`$manifests`

`$partitions`

`$files`



Trino + Iceberg Data App

Templated Queries

```
select * from <table_name>
```

Templated Dashboards

Runs a bunch of templated queries

Trino + Iceberg Data App

Iceberg Table Information

Shows the current production Iceberg table information ([go/iceberg-info](#))

 Leave feedback for @kevinliu



Schema and Partition Definition

Table Columns

Airflow Runs

Snapshot Info

Files Summary

History

Properties

Find table by s3 location

schema_name

table_name

Reset

Submit

Iceberg Schema and Partition Definition

Shows schema definition and partition definition

Create Table 

```
CREATE TABLE iceberg.fact.costs_raw (
```

Trino + Iceberg Data App

Iceberg Table Information ↑

Shows the current production Iceberg table information ([go/iceberg-info](#))

 Leave feedback for @kevinliu



Schema and Partition Definition

Table Columns

Airflow Runs

Snapshot Info

Files Summary

History

Properties

Find table by s3 location

schema_name

table_name

fact


costs_raw

Reset

Submit

Iceberg Airflow Writes

Associates Airflow execution_date to Iceberg snapshot id

'airflow_run (external-id)' 

committed_at 

snapshot_id 

spark_app_id 

operation 

search 8 records...

search 8 records...

search 8 records...

search 8 records...

search 8 records...

20230606

2023-06-07 11:41:38

6584378299722931065

application_1686087130904_0570

overwrite

20230605

2023-06-06 04:33:30

6697669884810570820

application_1685977767581_2025

overwrite

Trino + Iceberg Data App



Bonus features

- Analysis! Aggregation, filter, and search
- “When did the latest write happen?”
- “How much data was in the latest write? In terms of record count, file size.”
- “How big is my iceberg table?”
- “What is the schema and partition spec of my iceberg table?”
- “What is the average data file size?” For compaction.

Table Properties: The Missing Piece

- Iceberg stores metadata in the Hive Metastore
- “Retention Policy” and other stateful settings
- ✗ Not accessible by Trino

Inspect Table Metadata from Hive Metastore

Trino + PostgreSQL!

Trino PostgreSQL connector

Read Hive Metastore's database directly with Trino



Trino + PostgreSQL!

Iceberg Table Information ⌵

Shows the current production Iceberg table information ([go/iceberg-info](#))

🔊 Leave feedback for @kevinliu



Schema and Partition Definition

Table Columns

Airflow Runs

Snapshot Info

Files Summary

History

Properties

Find table by s3 location

schema_name

table_name

fact

costs_raw

Reset

Submit

Iceberg Properties

Properties set in the Hive Metastore. Data is read directly from the Hive Metastore DB

key ⌵

value ⌵

search 16 records...

search 16 records...

write.parquet.compression-codec

zstd

write.location-provider.impl

com.stripe.iceberg.LocalizedLocationProvider

Trino + PostgreSQL!



Bonus features

- Analysis! 10k foot view of our data warehouse
- “How many tables and views are there in our data warehouse?”
- “How many are hive tables versus iceberg tables?”
- “How many tables contain the `merchant_id` field?”

Success!



What about Iceberg Table Usage?

“I want to know how this table is used”

“I want to know who’s using this table”

“I want to deprecate this table and migrate to this new table”

(RUN DATASERVING-1628)

Inspecting Iceberg usage with Trino queries

Trino + PostgreSQL + Trino!

Trino event listener (`QueryCompletedEvent`)

Store events in PostgreSQL DB

Trino PostgreSQL connector to read the events



Query Metadata Enrichment

Modify Trino clients to send extra metadata using `X-Trino-Client-Info` in HTTP request

Example metadata:

- Dashboard ID
- Airflow task name
- Checker ID
- Workflow engine job name

Query Metadata Enrichment

Show up in the source table as `client_info` field

JSON allows arbitrary metadata encoded as dictionary

Access as dictionary using Trino

```
json_extract_scalar(client_info, '$.host')
```

Query Metadata Enrichment



Bonus features

- **Analysis!** Enables multitude of auditing, debugging, and optimization use cases
- “Who is still using this table? I need to deprecate it!”
- “What are the most queried tables?”
- “What are the most queried columns for this table?”
- “What is the top service using Trino? What is the usage pattern?”
- “What are the most common errors in the last hour?”

Future Work!

Trino as validation framework

Trino for Iceberg table maintenance (compaction, manifest rewrite)

Table optimizer based on read patterns

Hive Metastore Attribution

Q & A

stripe



Thank you!

 kevinjqliu

 kevinjqliu