



Insert / Update / Delete on S3 With Amazon Athena and Apache Iceberg

Ramkumar Nottath
Sr. Solutions Architect (Analytics)



Insert / Update / Delete on S3 With Amazon Athena and Apache Iceberg

Ramkumar Nottath
Sr. Solutions Architect (Analytics)

Amazon Athena

QUERY ALL YOUR
DATA USING SQL



Serverless



Open and flexible



Fast and interactive



Cost effective

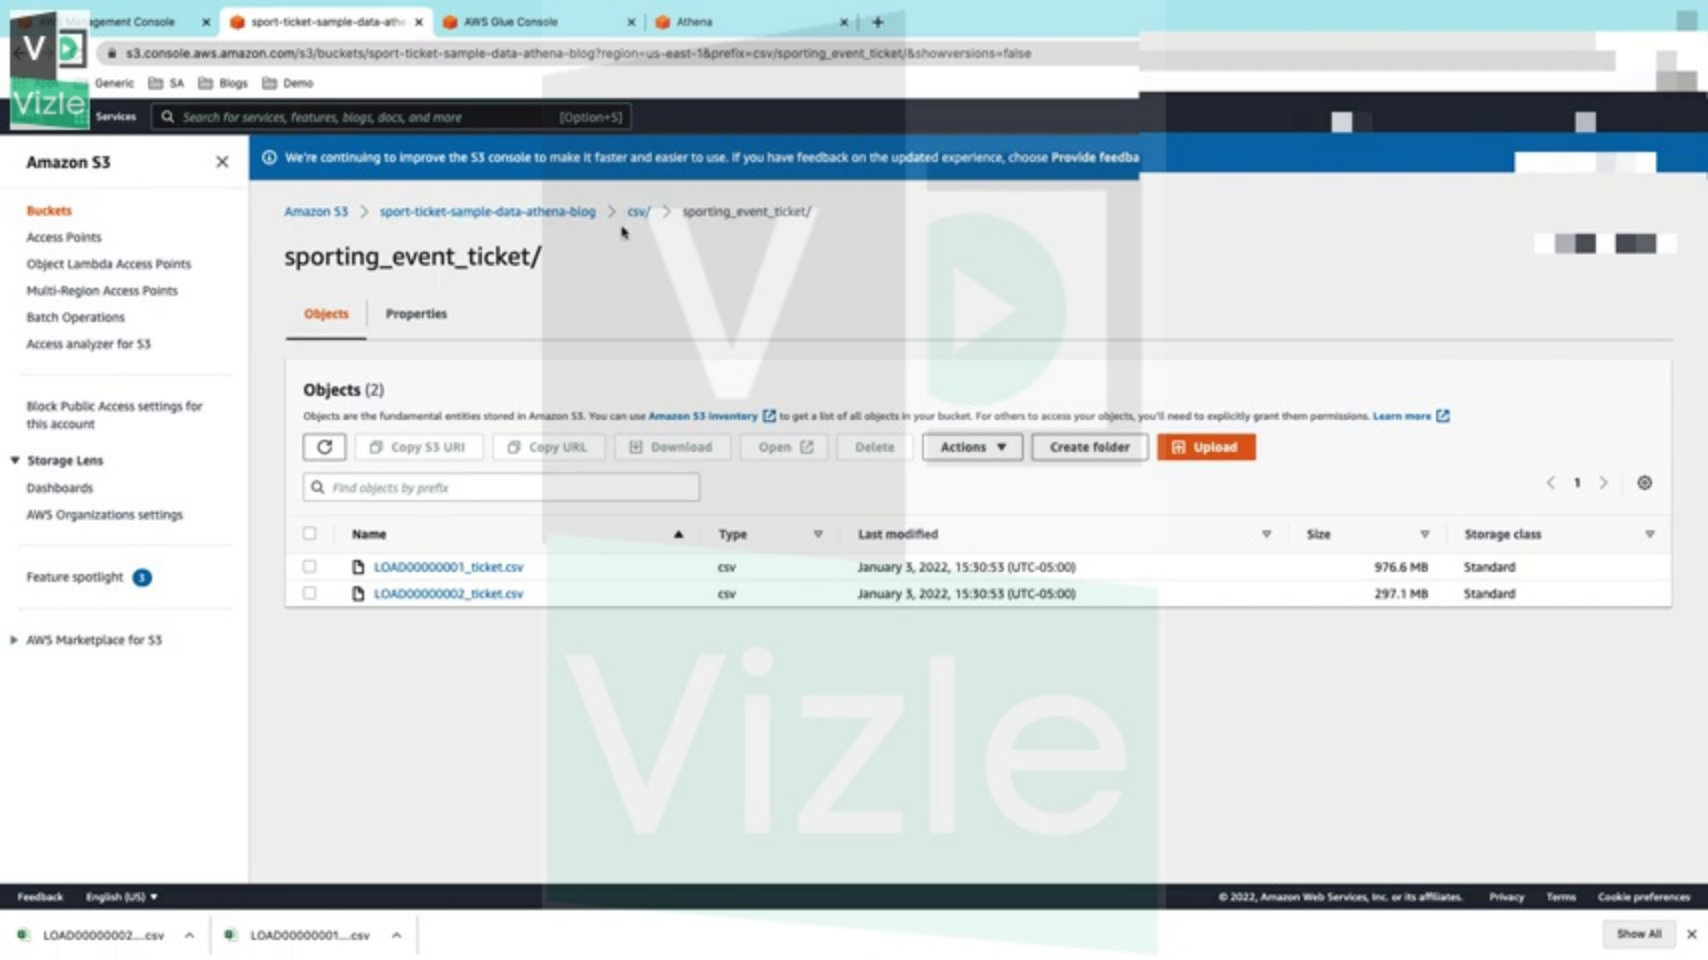


Athena ACID Transactions powered by Apache Iceberg

OPEN SPEC, SNAPSHOT-BASED TABLE
FORMAT FOR HUGE ANALYTIC DATASETS



ICEBERG



Amazon S3

- Buckets
- Access Points
- Object Lambda Access Points
- Multi-Region Access Points
- Batch Operations
- Access analyzer for S3

Block Public Access settings for this account

- Storage Lens
- Dashboards
- AWS Organizations settings

Feature spotlight

AWS Marketplace for S3

We're continuing to improve the S3 console to make it faster and easier to use. If you have feedback on the updated experience, choose Provide feedback

Amazon S3 > sport-ticket-sample-data-athena-blog > csv/ > sporting_event_ticket/

sporting_event_ticket/

Objects Properties

Objects (2)

Objects are the fundamental entities stored in Amazon S3. You can use [Amazon S3 Inventory](#) to get a list of all objects in your bucket. For others to access your objects, you'll need to explicitly grant them permissions. [Learn more](#)

- Refresh
- Copy S3 URI
- Copy URL
- Download
- Open
- Delete
- Actions
- Create folder
- Upload

Find objects by prefix

<input type="checkbox"/>	Name	Type	Last modified	Size	Storage class
<input type="checkbox"/>	LOAD000000001_ticket.csv	csv	January 3, 2022, 15:30:53 (UTC-05:00)	976.6 MB	Standard
<input type="checkbox"/>	LOAD000000002_ticket.csv	csv	January 3, 2022, 15:30:53 (UTC-05:00)	297.1 MB	Standard

AWS Glue Edit table Delete table

Data catalog
 Databases
 Tables
 Connections
 Crawlers
 Classifiers
 Schema registries
 Schemas
 Settings

ETL
 AWS Glue Studio
 Blueprints
 Workflows
 Jobs
 ML Transforms
 Triggers
 Dev endpoints
 Notebooks
 Security
 Security configurations
 Tutorials
 Add crawler

Name sporting_event_ticket
Description ticketdata
Database ticketdata
Classification csv
Location s3://sport-ticket-sample-data-athena-biog/csv/sporting_event_ticket/
Connection
Deprecated No
Last updated Mon Jan 03 16:22:14 GMT-500 2022
Input format org.apache.hadoop.mapred.TextInputFormat
Output format org.apache.hadoop.hive.q1.io.HiveIgnoreKeyTextOutputFormat
Serde serialization lib org.apache.hadoop.hive.serde2.lazy.LazySimpleSerDe
Serde parameters field.delim ,

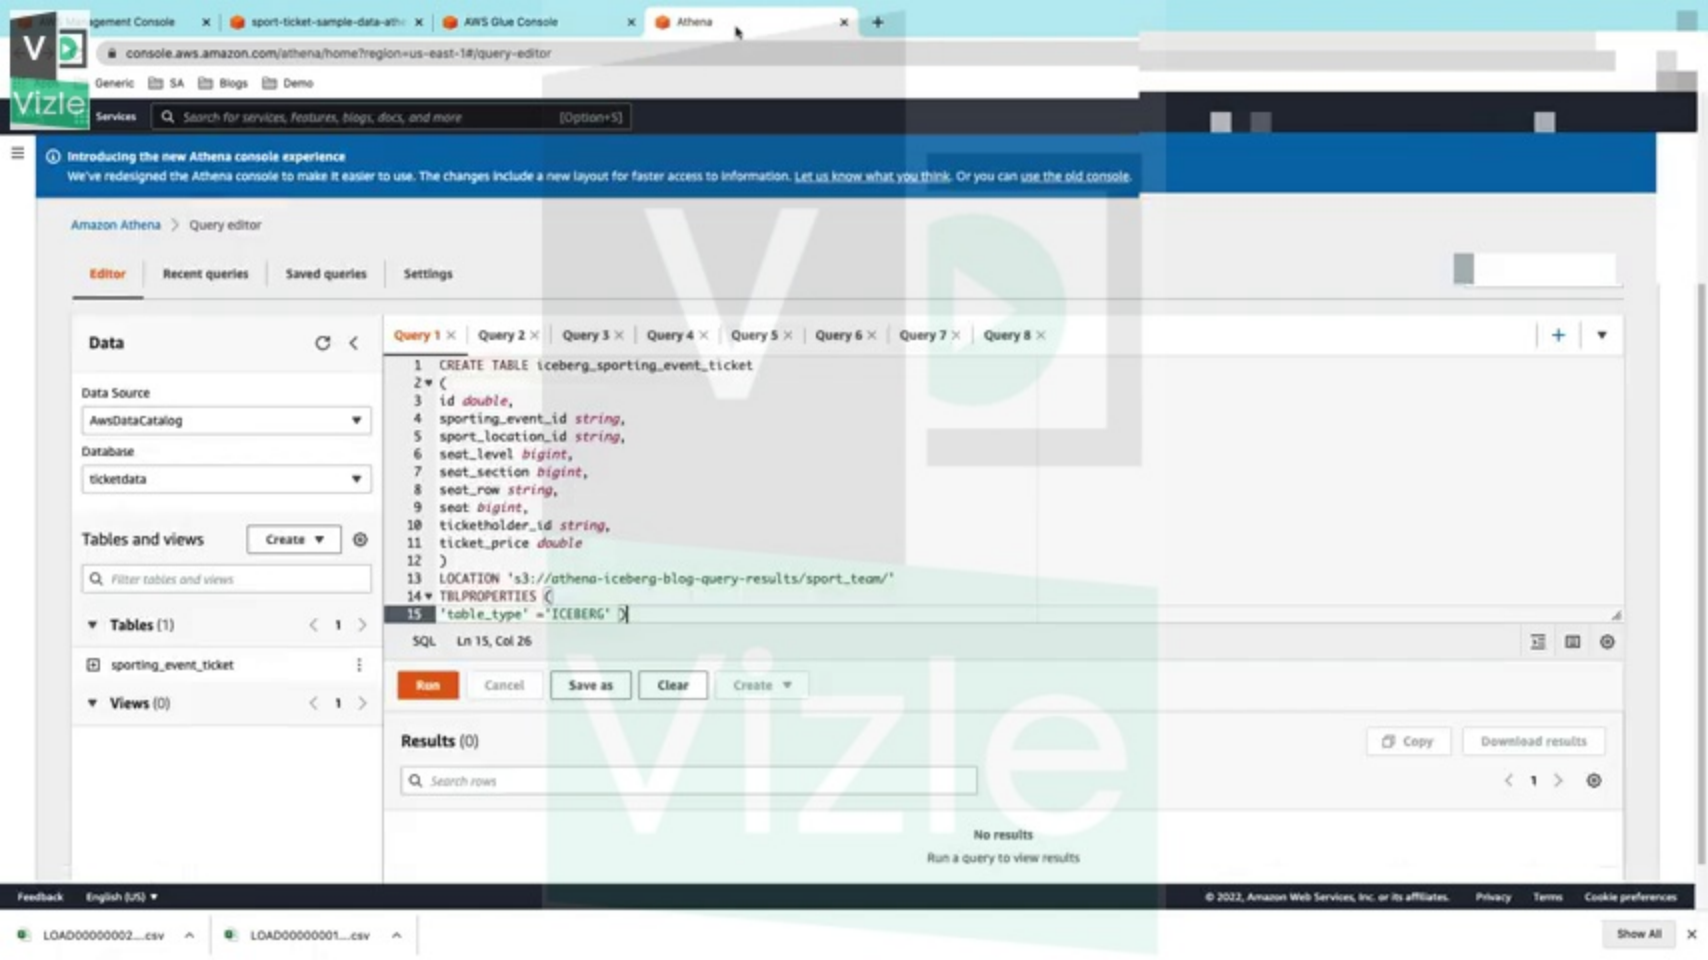
Table properties

skip.header.line.count	1	sizeKey	1335544685	objectCount	2	UPDATED_BY_CRAWLER	ticket-crawler	CrawlerSchemaSerializerVersion	1.0	recordCount	14360695		
averageRecordSize	93	CrawlerSchemaDeserializerVersion	1.0	compressionType	none	columnsOrdered	true	areColumnsQuoted	false	delimiter	,	typeOfData	file

Schema

Showing: 1 - 9 of 9

	Column name	Data type	Partition key	Comment
1	id	double		
2	sporting_event_id	string		
3	sport_location_id	string		
4	seat_level	bigint		
5	seat_section	bigint		
6	seat_row	string		



Introducing the new Athena console experience
We've redesigned the Athena console to make it easier to use. The changes include a new layout for faster access to information. Let us know what you think. Or you can use the old console.

Amazon Athena > Query editor

Editor Recent queries Saved queries Settings

Data

Data Source
AwsDataCatalog

Database
ticketdata

Tables and views Create

Filter tables and views

Tables (1) < 1 >

sporting_event_ticket

Views (0) < 1 >

Query 1 x Query 2 x Query 3 x Query 4 x Query 5 x Query 6 x Query 7 x Query 8 x

```
1 CREATE TABLE iceberg_sporting_event_ticket
2 (
3   id double,
4   sporting_event_id string,
5   sport_location_id string,
6   seat_level bigint,
7   seat_section bigint,
8   seat_row string,
9   seat bigint,
10  ticketholder_id string,
11  ticket_price double
12 )
13 LOCATION 's3://athena-iceberg-blog-query-results/sport_team/'
14 TBLPROPERTIES (
15   'table_type' = 'ICEBERG' )
```

SQL Ln 15, Col 26

Run Cancel Save as Clear Create

Results (0) Copy Download results

Search rows

No results
Run a query to view results

Data

Data Source

AwsDataCatalog

Database

ticketdata

Tables and views

Create

Filter tables and views

Tables (2)

iceberg_sporting_event_ticket

sporting_event_ticket

Views (0)

Query 1 x Query 2 x Query 3 x Query 4 x Query 5 x Query 6 x Query 7 x Query 8 x

```
1 select count(*) from ticketdata.sporting_event_ticket;
```

SQL Ln 1, Col 40

Run again

Cancel

Save as

Clear

Create

Completed

Time in queue: 0.587 sec

Run time: 1.055 sec

Data scanned: 1.24 GB

Results (1)

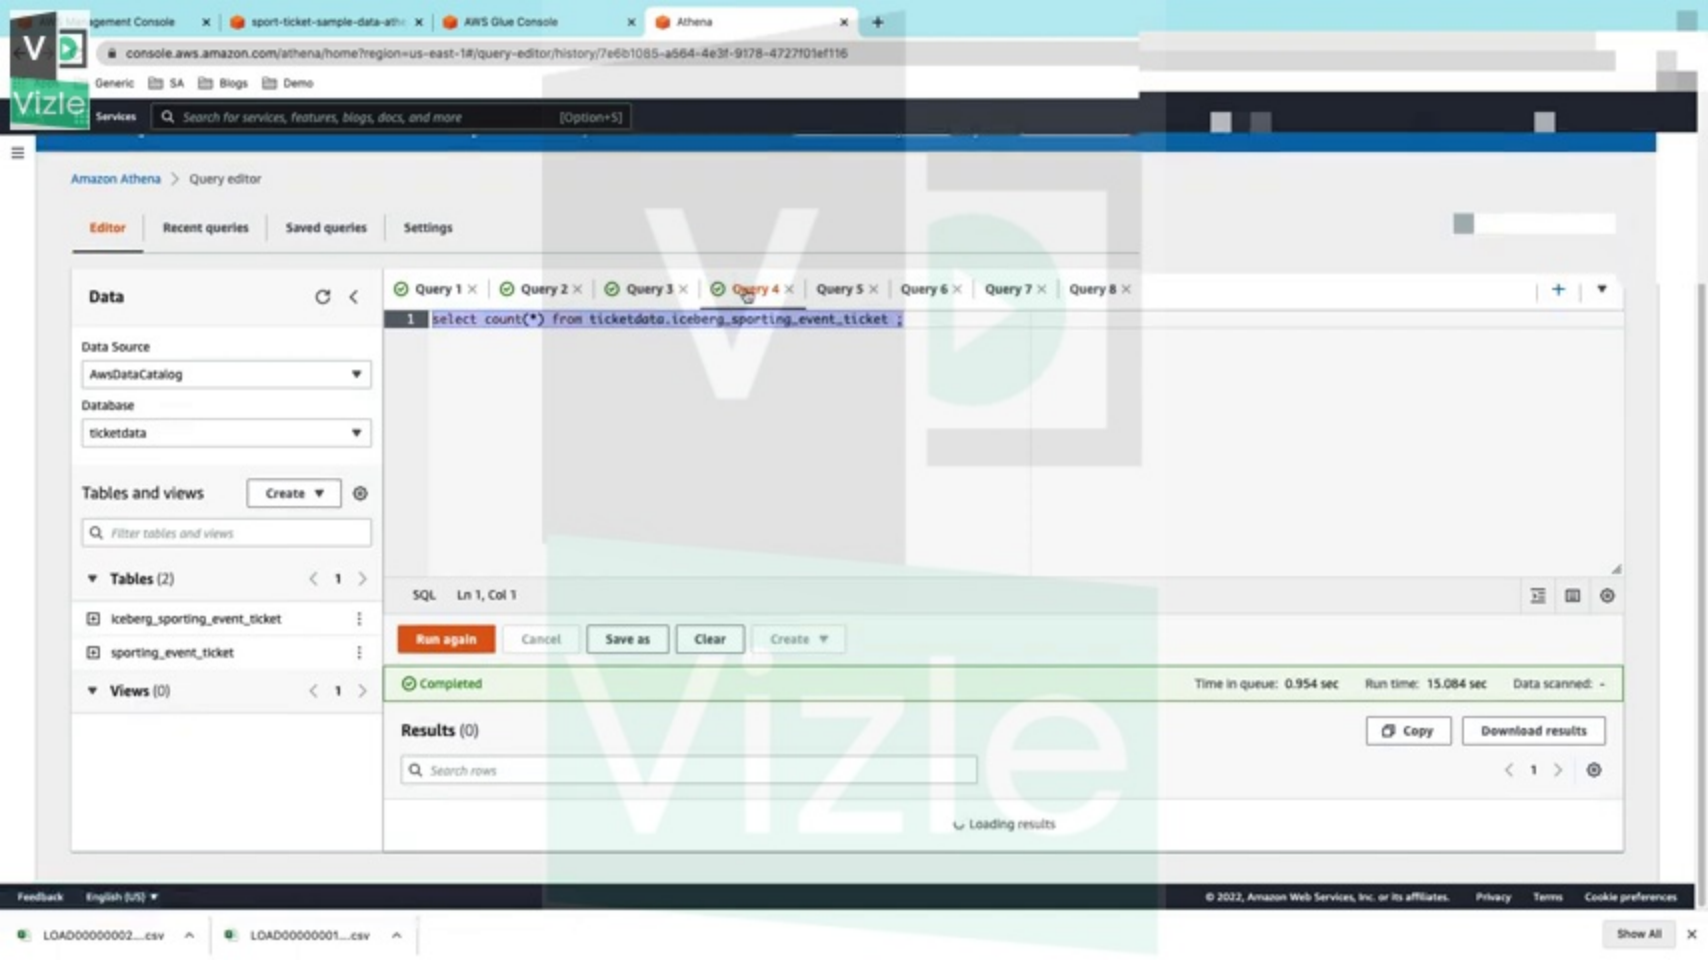
Copy

Download results

Search rows

_col0

14997376



Data

Data Source

AwsDataCatalog

Database

ticketdata

Tables and views

Create

Filter tables and views

Tables (2)

iceberg_sporting_event_ticket

sporting_event_ticket

Views (0)

Query 1 x Query 2 x Query 3 x Query 4 x Query 5 x Query 6 x Query 7 x Query 8 x

```
1 select count(*) from ticketdata.iceberg_sporting_event_ticket;
```

SQL Ln 1, Col 1

Run again

Cancel

Save as

Clear

Create

Completed

Time in queue: 0.954 sec

Run time: 15.084 sec

Data scanned: -

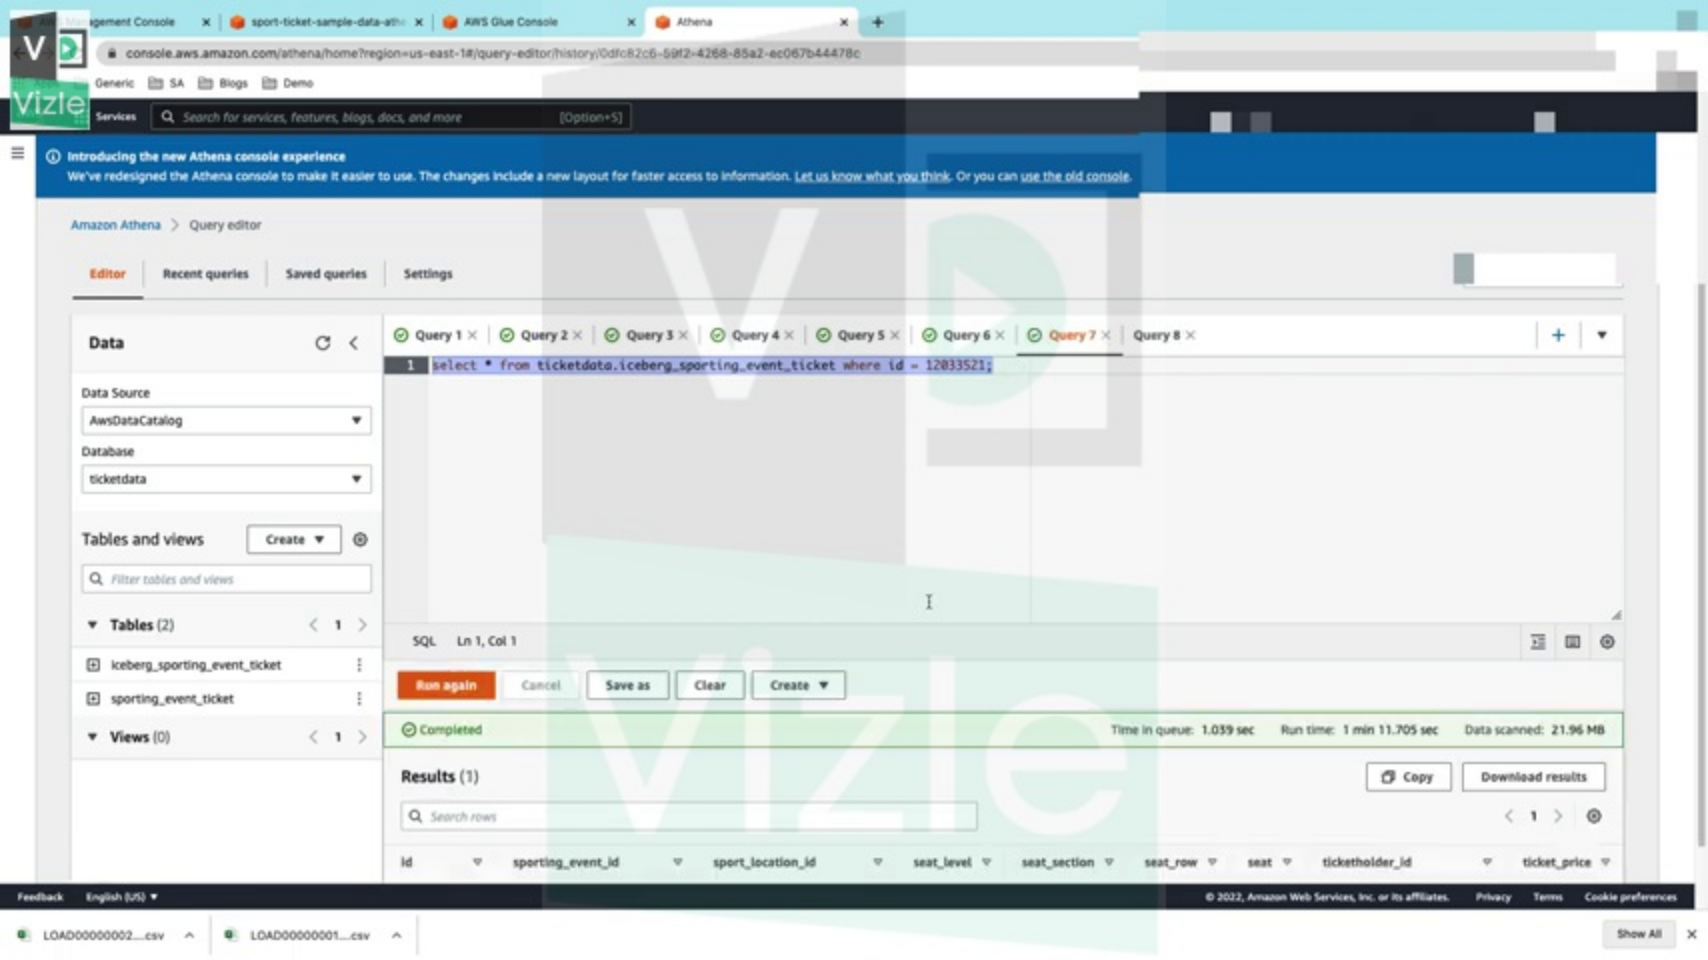
Results (0)

Copy

Download results

Search rows

Loading results



Introducing the new Athena console experience. We've redesigned the Athena console to make it easier to use. The changes include a new layout for faster access to information. Let us know what you think. Or you can use the old console.

Amazon Athena > Query editor

Editor Recent queries Saved queries Settings

Data

Data Source: AwsDataCatalog Database: ticketdata

- Tables and views: Create Filter tables and views Tables (2): iceberg_sporting_event_ticket, sporting_event_ticket Views (0)

Query 1 x Query 2 x Query 3 x Query 4 x Query 5 x Query 6 x Query 7 x Query 8 x

```
1 select * from ticketdata.iceberg_sporting_event_ticket where id = 12833521;
```

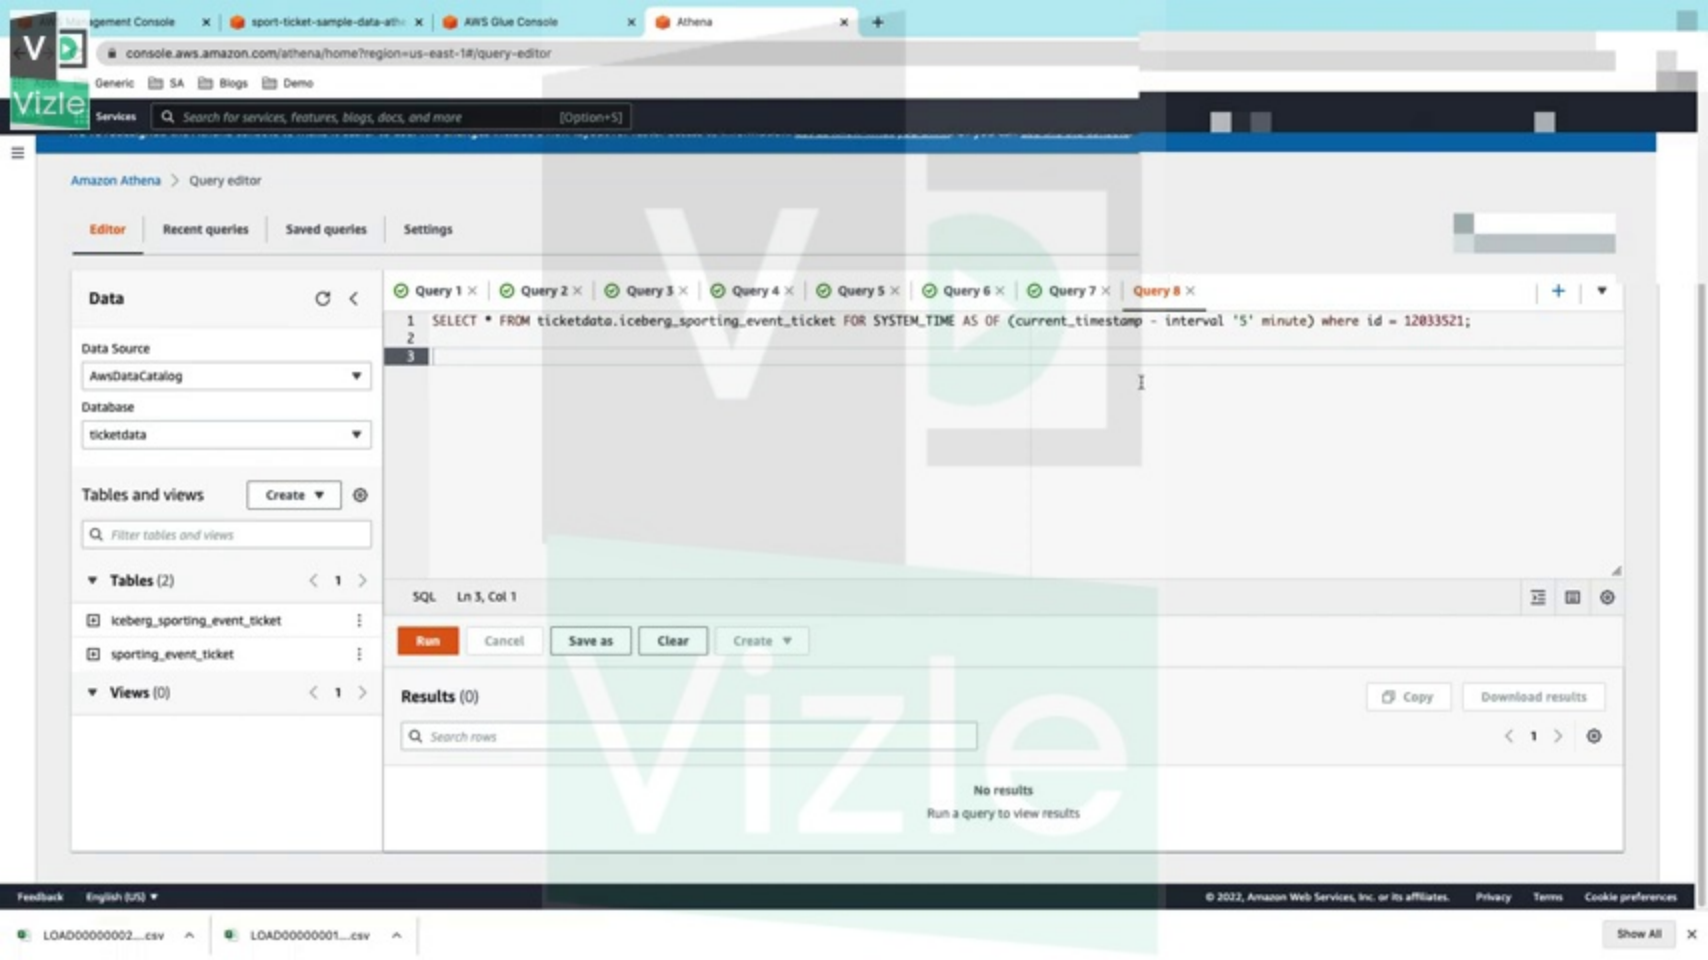
SQL Ln 1, Col 1 Run again Cancel Save as Clear Create

Completed Time in queue: 1.039 sec Run time: 1 min 11.705 sec Data scanned: 21.96 MB

Results (1) Copy Download results

Search rows

Table header: id, sporting_event_id, sport_location_id, seat_level, seat_section, seat_row, seat, ticketholder_id, ticket_price



Data

Data Source

AwsDataCatalog

Database

ticketdata

Tables and views

Create

Filter tables and views

Tables (2)

iceberg_sporting_event_ticket

sporting_event_ticket

Views (0)

Query 1 x Query 2 x Query 3 x Query 4 x Query 5 x Query 6 x Query 7 x Query 8 x

```
1 SELECT * FROM ticketdata.iceberg_sporting_event_ticket FOR SYSTEM_TIME AS OF (current_timestamp - interval '5' minute) where id = 12033521;
2
3
```

SQL Ln 3, Col 1

Run

Cancel

Save as

Clear

Create

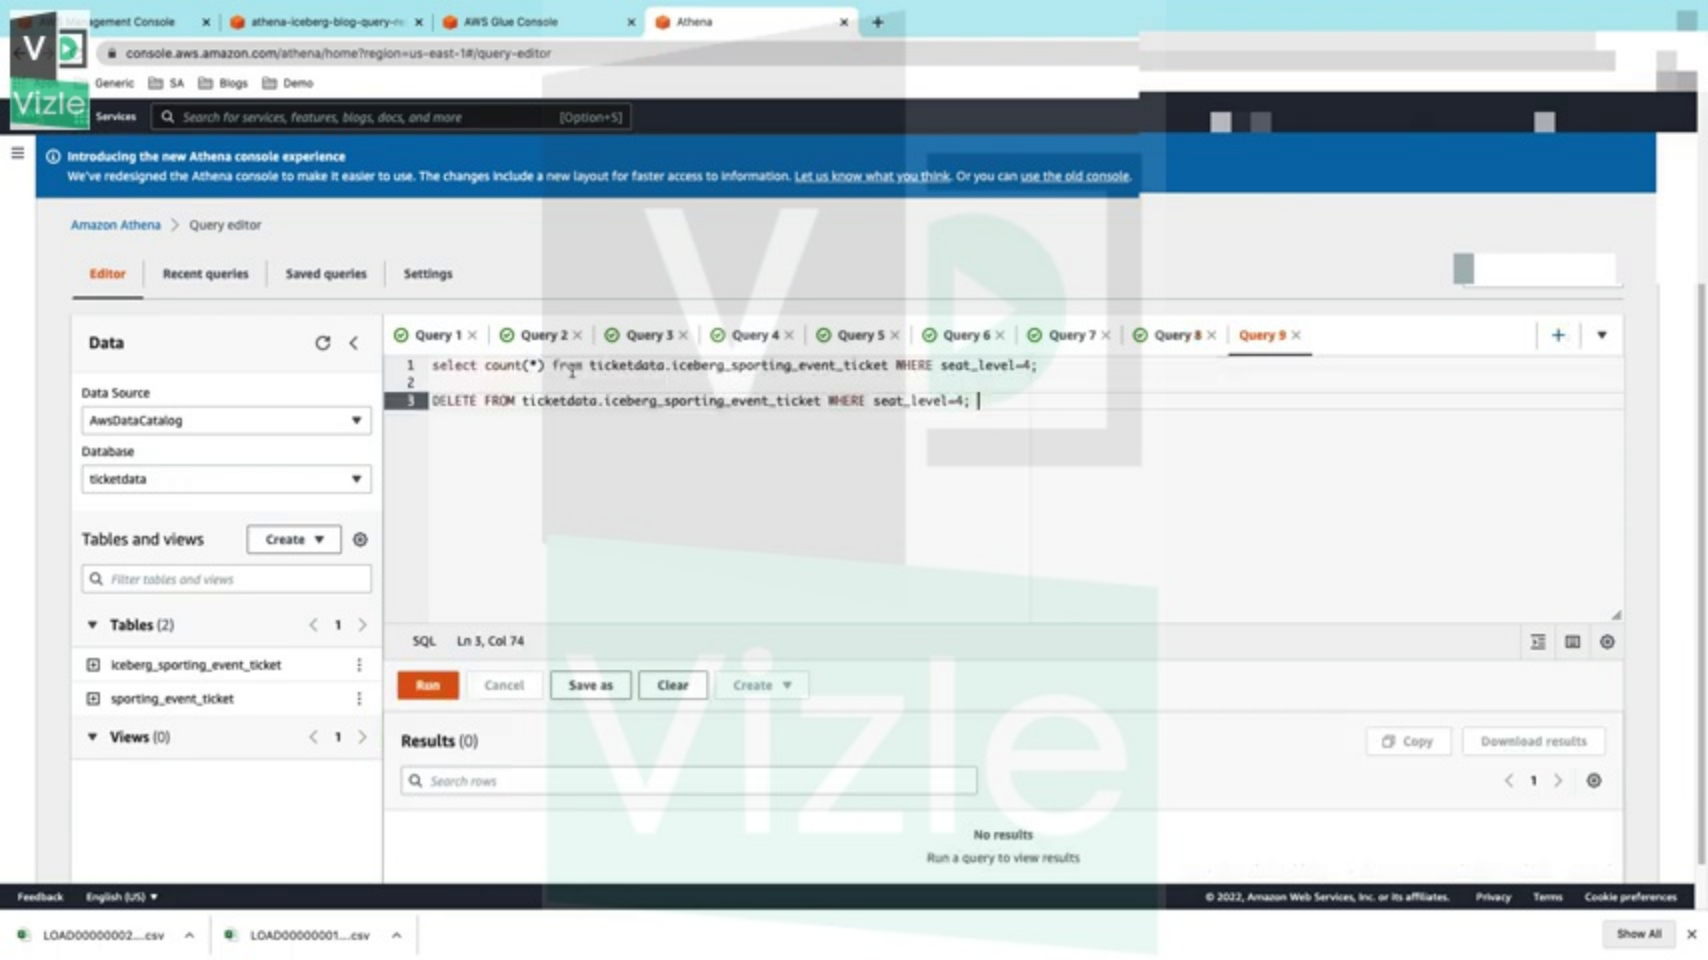
Results (0)

Search rows

Copy

Download results

No results
Run a query to view results



Introducing the new Athena console experience
We've redesigned the Athena console to make it easier to use. The changes include a new layout for faster access to information. Let us know what you think. Or you can use the old console.

Amazon Athena > Query editor

Editor Recent queries Saved queries Settings

Data

Data Source
AwsDataCatalog
Database
ticketdata

Tables and views Create
Filter tables and views
Tables (2)
iceberg_sporting_event_ticket
sporting_event_ticket
Views (0)

Query 1 x Query 2 x Query 3 x Query 4 x Query 5 x Query 6 x Query 7 x Query 8 x Query 9 x

```
1 select count(*) from ticketdata.iceberg_sporting_event_ticket WHERE seat_level=4;
2
3 DELETE FROM ticketdata.iceberg_sporting_event_ticket WHERE seat_level=4;
```

SQL Ln 3, Col 74
Run Cancel Save as Clear Create

Results (0) Copy Download results
Search rows

No results
Run a query to view results

Data

Data Source

AwsDataCatalog

Database

ticketdata

Tables and views

Create

Filter tables and views

Tables (2)

iceberg_sporting_event_ticket

sporting_event_ticket

Views (0)

Query 1 x Query 2 x Query 3 x Query 4 x Query 5 x Query 6 x Query 7 x
1 select count(*) from ticketdata.iceberg_sporting_event_ticket WHERE seat_level=4;
2
3 DELETE FROM ticketdata.iceberg_sporting_event_ticket WHERE seat_level=4;

SQL Ln 1, Col 83

Run again

Cancel

Save as

Clear

Create

Completed

Time in queue: 1.458 sec

Run time: 17.015 sec

Data scanned: -

Results (1)

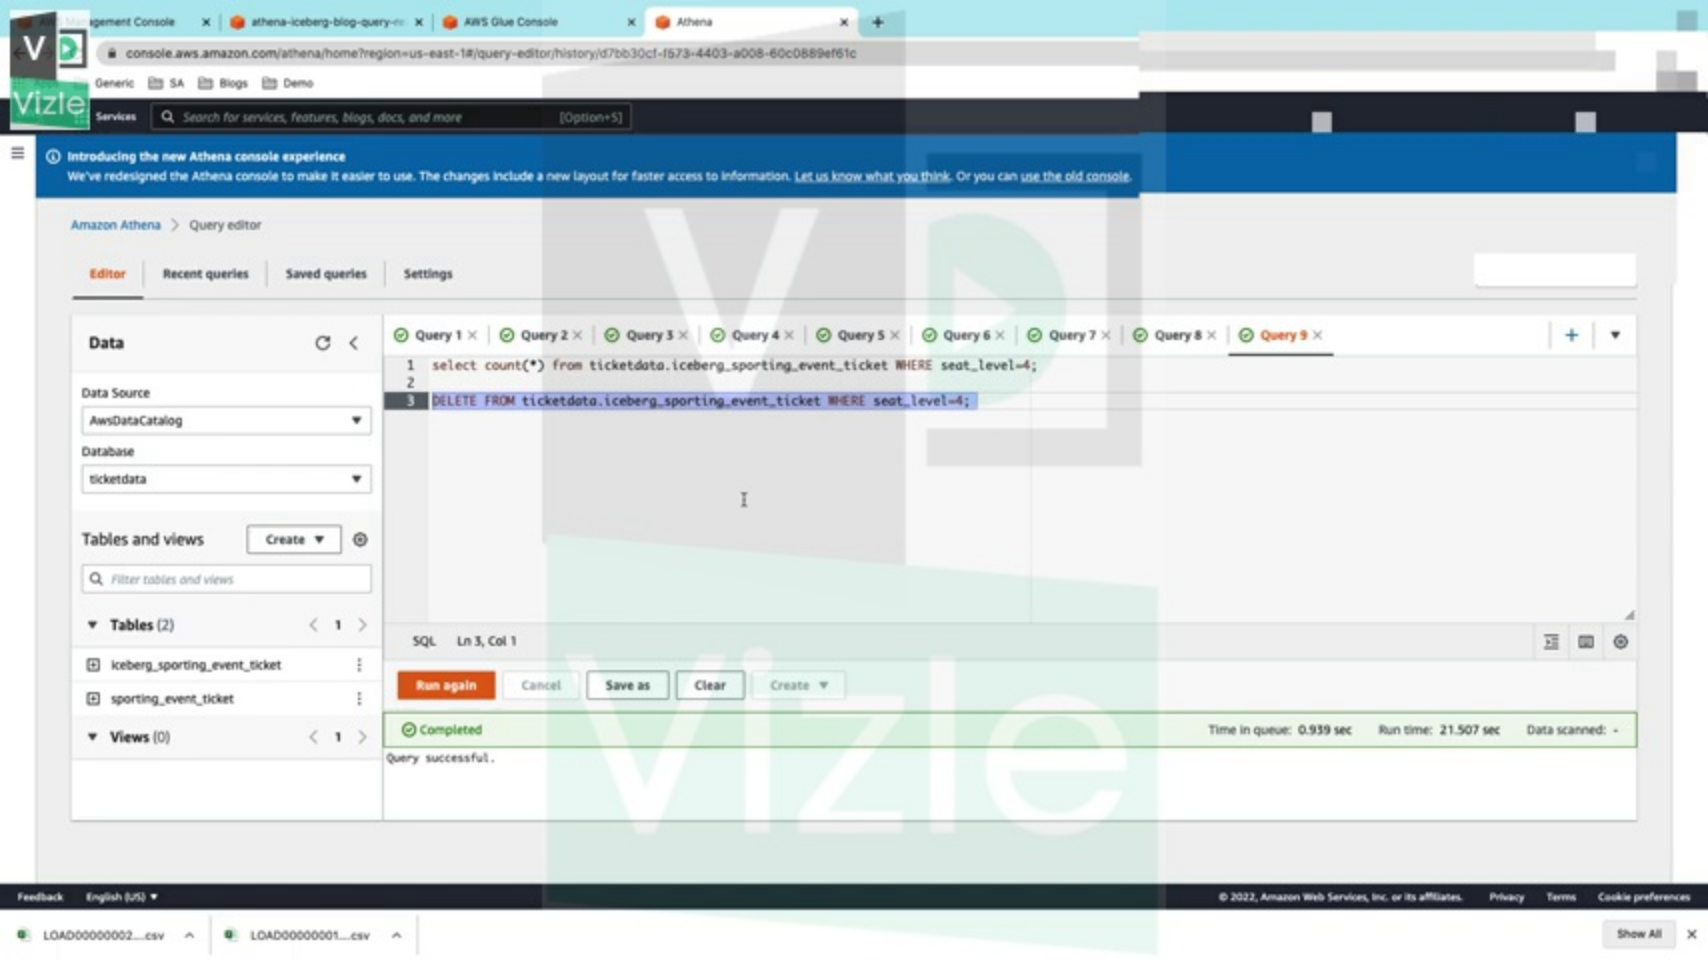
Copy

Download results

Search rows

_col0

140j873



Introducing the new Athena console experience

We've redesigned the Athena console to make it easier to use. The changes include a new layout for faster access to information. Let us know what you think. Or you can use the old console.

Amazon Athena > Query editor

Editor Recent queries Saved queries Settings

Data

Data Source

AwsDataCatalog

Database

ticketdata

Tables and views

Create

Filter tables and views

Tables (2)

iceberg_sporting_event_ticket

sporting_event_ticket

Views (0)

Query 1 x Query 2 x Query 3 x Query 4 x Query 5 x Query 6 x Query 7 x Query 8 x Query 9 x

```
1 select count(*) from ticketdata.iceberg_sporting_event_ticket WHERE seat_level=4;
2
3 DELETE FROM ticketdata.iceberg_sporting_event_ticket WHERE seat_level=4;
```

SQL Ln 3, Col 1

Run again

Cancel

Save as

Clear

Create

Completed

Time in queue: 0.939 sec

Run time: 21.507 sec

Data scanned: -

Query successful.

Data

Data Source

AwsDataCatalog

Database

ticketdata

Tables and views

Create

Filter tables and views

Tables (2)

iceberg_sporting_event_ticket

sporting_event_ticket

Views (0)

Query 1 x Query 2 x Query 3 x Query 4 x Query 5 x Query 6 x Query 7 x
1 select count(*) from ticketdata.iceberg_sporting_event_ticket WHERE seat_level=4;
2
3 DELETE FROM ticketdata.iceberg_sporting_event_ticket WHERE seat_level=4;

SQL Ln 1, Col 1

Run again

Cancel

Save as

Clear

Create

Completed

Time in queue: 0.889 sec

Run time: 25.807 sec

Data scanned: -

Results (1)

Copy

Download results

Search rows

_col0

0



What you learned today

- Creating ICEBERG table in Athena
- Insert, Update, Delete operations on Amazon S3
- Time Travel query

This PDF is generated automatically by **Vizle**.
Slides created *only for a few minutes* of your Video.



For the full PDF, please **Login to Vizle**.

<https://vizle.offnote.co> (Login via Google, top-right)

Stay connected with us:

Join us on **Facebook, Discord, Quora, Telegram**.