

Rec Room 🤝 RudderStack

How Rec Room Processes Billions of Events Per Day with Databricks and RudderStack

Databricks 2023







Albert Hu

Senior Analytics Engineer Rec Room



Lewis Mbae

Head of Customer Engineering RudderStack



About Rec Room



- Rec Room is the best place to build and play games together.
- With more than 80M lifetime users, you can party up with friends from all around the world to play, hang out, explore MILLIONS of player-created rooms, or build something new!
- Founded in 2016



More Users, More Data



Making Impact with Data at Rec Room

Data-driven culture powered by high data volume and quick adoption

Ц

A/B Test as many decisions and features as possible ◆

Recommend new, interesting rooms and items to players iii

Share metrics with Creators to help grow their player base

Rec Room's Data Challenges

Disparate data sources

Disparate data destinations

Multiple environments



.NET

microsoft

JS

javascript

unity

flutter

webhook

150+ experiments launched as of Q2 2023

Started A/B testing March 2022

One A/B test, big impact

From a negative to positive 25% change in chat messages between players





About RudderStack



- RudderStack is the leading lakehouse native customer data platform.
- RudderStack runs on top of your lakehouse and does not store data, alleviating security concerns, reducing costs, and unlocking the value of your lakehouse investment
- Founded in 2019





Rec Room Data Architecture



What is RudderStack Transformations?

Benefits:



How Rec Room uses Transformations



Filtering events by destination



Cleaning and enriching event data before it lands in the destination



Testing Transformations in dev environments before shipping to production

Tips for RudderStack Transformations

Treat RudderStack Transformations like any other code

C Define functions that are used repeatedly to keep code DRY





Automate testing using unit tests via Github Actions

Simplifying Databrick's Data Processing

RudderStack's clean lakehouse data model enables painless data processing

identities Table		users Table		tracks Table	
user_id:_	varchar	user_id:_	varchar	user_id:_	varchar
anonymous_id:_	varchar	email:_	varchar	anonymous_id:_	varchar
email:_	varchar	first_name:_	varchar	event:_	varchar
first_name:_	varchar	last_name:_	varchar	context_ <prop>:_</prop>	varchar
last_name:_	varchar	context_ <prop>:_</prop>	varchar	timestamp:_	timestamp
context_ <prop>:_</prop>	varchar	timestamp:_	timestamp		
timestamp:_	timestamp				

Simplifying Databricks' Data Processing

Using a medallion data architecture to power multiple use-cases

Gold Silver Bronze Filtered, cleaned, augmented **Raw ingestion**

- Copied from blob storage and • inserted hourly
- No table stats collection

- Merge w/ 2-day lookback window to process late arriving data
- Partition by server received date
- Table stats collected Daily

Business-facing

- Append by server received date
- Partition by event date •
- Table stats collected Weekly

Configure tblproperties sooner than later

 Configure appropriate
 tblproperties during the pipeline building process

 Structure table columns to optimize z-ordering and take advantage of data skipping by data type: keys/numericals on the left and strings on the right



ALTER TABLE example_table SET TBLPROPERTIES ('delta.columnMapping.mode' = 'name', 'delta.minReaderVersion' = '2', 'delta.minWriterVersion' = '5', 'delta.dataSkippingNumIndexedCols' = 40);

Choose the right compute

Choose the appropriate compute based on **frequency** and **processing time (including cluster startup)**

- For **Scheduled automated jobs** check if **Job** vs. **All-Purpose Cluster** are appropriate (3x cost difference)
- For SQL-only job , try Serverless SQL Warehouses

Cluster * 🕐	Job_cluster 126 GB · 36 Cores · DBR 12.2 LTS · Spark 3.3.2 · Scala 2.12 🖉 💊	-
Dependent libraries 🕐	Job Clusters	
Parameters ?	Job_cluster 126 GB · 36 Cores · DBR 12.2 LTS · Spark 3.3.2 · Scala 2.12	9
Emails ③	Add new job cluster	
Retries 🕐	Existing All-Purpose Clusters	2
Timeout in seconds 🕐	DBR 12.0 ML · Spark 3.3.1 · Scala 2.12	
	UnityCatalogShared	3

Optimize compute cost over time

Map workloads to compute that optimizes **cost over time**

 Selecting the smallest compute is not always cheapest (and vice versa)

 Find clues workloads might not be compatible with compute such as if bytes are spilling to disk

Aggregated task time U	
Tasks total time	67.34 h
Tasks time in Photon	61 %
ю	
Rows read	4,945,103,885
Bytes read	51.33 GE
Bytes read from cache	0 %
Bytes written	0 by
Files & partitions	
Files read	2,405
Partitions read	2,192
Spilling	
Bytes spilled to disk	4.84 TE

Choosing a partition

Query the **DeltaLog** to understand if a table is partitioned appropriately

- Table Size >1TB
- Size of each partition is **1GB+**
- Field is low cardinality and will be used for filtering and merge operations

	vac parcitions = me	eca.				
2	.select(explode("partitionValue	is"))			
3	.select(\$"key")					
4	.distinct					
5	.map(f=>f.getStr	ing(0))				
6	.collect()					
7	.toList					
8	10 10 10 10 10 10 10 10 10 10 10 10 10 1					
9	val partitionsClaus	se = partitions.	<pre>map(x => col(s"parti1</pre>	tion	Values.\${x}"))	
10						
11	display(`			
12	meta.groupBy(part	<pre>L) co(loi=cluse:_*</pre>)			
13	.agg(Sull("S12e	·).ds(~size~))	nell/lit/hutee to mb d	Faat	((ma)	
14	.wrthcotumn("s	12e_111_000, \$~\$1	ze"/tit(bytes_to_mb_i	act	.01))	
15	.withColumn("s	ize in gb". S"si	ze"/lit(bytes to gb f	Fact	or))	
15 16	<pre>.withColumn("s')</pre>	ize_in_gb", \$"si	<pre>ze"/lit(bytes_to_gb_1</pre>	fact	or))	
15 16 (5) Tab	.withColumn("s")) Spark Jobs le v +	ize_in_gb", \$"si	<pre>ze"/lit(bytes_to_gb_1</pre>	fact	or))	
15 16 ▶ (5) Tab	.withColumn("s")) Spark Jobs de v + created_date_utc v	ize_in_gb", \$"si size ▲	<pre>ze"/lit(bytes_to_gb_1 size_in_mb</pre>	fact	or)) size_in_gb	4
15 16 ▶ (5) Tab	y spark Jobs le v + created_date_utc ▼ 2023-06-04	ize_in_gb", \$"si size ▲ 23224942825	<pre>ze"/lit(bytes_to_gb_1 size_in_mb 22149.031472206116</pre>	Fact ▲	or)) size_in_gb 21.629913547076285	4
15 16 (5) Tab 1 2	.withColumn("s") Spark Jobs e	ize_in_gb", \$"si size ▲ 23224942825 23141441690	ze"/lit(bytes_to_gb_1 size_in_mb 22149.031472206116 22069.398584365845	Fact ▲	size_in_gb 21.629913547076285 21.55214705504477	4
15 16 (5) Tab 1 2 3	.withColumn("s:) Spark Jobs ile ∨ + created_date_utc ♥ 2023-06-04 2023-06-03 2023-06-02	<pre>size</pre>	<pre>ze"/lit(bytes_to_gb_1 size_in_mb 22149.031472206116 22069.398584365845 19133.15301799774</pre>	▲	size_in_gb 21.629913547076285 21.55214705504477 18.68471974413842	4
15 16 (5) Tab 1 2 3 4	.withColumn("s") Spark Jobs le ↓ created_date_utc ↓ 2023-06-04 2023-06-03 2023-06-02 2023-06-02 2023-06-01 1	size 23224942825 23141441690 20062565059 16708152819	<pre>ze"/lit(bytes_to_gb_1 size_in_mb 22149.031472206116 22069.398584365845 19133.15301799774 15934.136218070984</pre>	▲	size_in_gb 21.629913647076285 21.55214705504477 18.68471974413842 15.560679900459945	4
15 16 (5) Tab 1 2 3 4 5	.withColumn("s") Spark Jobs le ∨ + created_date_utc ♥ 2023-06-04 2023-06-03 2023-06-02 2023-06-01 2023-05-31	size 23224942825 23141441690 20062565059 16708152819 14776246804	<pre>size_in_mb 22149.031472206116 22069.398584365845 19133.15301799774 15934.136218070984 14091.727069854736</pre>	▲	size_in_gb 21.629913547076285 21.55214705504477 18.6847197413842 15.560679900459945 13.761452216655016	4
15 16 ► (5) Tab 1 2 3 4 5 6	.withColumn("s") Spark Jobs le ✓ + created_date_utc ✓ 2023-06-04 2023-06-03 2023-06-02 2023-06-01 2023-06-31 2023-05-30	size 23224942825 23141441690 20062565059 16708152819 14776246804 14721323634	<pre>ze"/lit(bytes_to_gb_1 size_in_mb 22149.031472206116 22069.398584365845 19133.15301799774 15934.136218070984 14091.727069854736 14039.348253250122</pre>	▲	size_in_gb 21.629913547076285 21.55214705504477 18.68471974413842 15.560679900459945 13.761452216655016 13.710301028564572	4

Prune Partitions to Optimize Merge

Rewrite the least number of files

Filter on the partition during merge





3x time difference after partition pruning

Maximize Photon Runtime

If you have Photon enabled, maximize the % of tasks using Photon

- Query performance can be dramatically faster on Photon (2x+)
- Check if queries are using
 Photon-compatible functions

```
.
select
   user id
    , anonymous_id
from
    core.events.unified_events
where
    created date utc >= DATE SUB(current date, 100)
qualify row_number() over (partition by user_id order by
created at utc desc) = 1
limit 1000
select
    user id
     max_by(anonymous_id, created_at_utc) as anonymous_id
from
    core.events.unified_events
where
    created_date_utc >= date_sub(current_date, 100)
group by 1
limit 1000
```

Measure and Meet

- Analyze job metadata and create dashboards to measure pipeline processing time. **Translate time into dollars**
- Create team rituals to review and understand drivers behind trends



RudderStack + Databricks: Unlocking value benefits for Rec Room

Clean, high-quality data

- With RudderStack, we can collect data once and unify it across our entire stack
- Rudderstack is the backbone of A/B testing the user experience, providing us product insights as well as additional data on user preferences

Seamless integrations

- Engineering teams are spending more time on product features and less time troubleshooting integrations
- Transformations enable the team to quickly make changes to existing events

Unlocking business outcomes

- A/B testing (**150+ experiments** executed through Statsig as of Q2 2023)
- Clean data for product analytics in Amplitude
- Recommendation models to improve Discovery

What's Next for RecRoom?

- Enable even more **personalized** recommendations for our players
- Share new metrics and datasets with our Creator Community
- Create tools to automate the optimization decisions needed to scale the lakehouse



What's Next for RudderStack?

- Build unified 360 view of a customer on top of the lakehouse to power more impactful use-cases
- Bolster data governance offering for collection of high quality first-party data
- Provide **more out-of-the box integrations** for our customers
- Unlock **real-time use cases** with near real-time sync to Databricks

rudderstack

The Lakehouse Native CDP

Thank you for Attending!

Albert Hu

ahu@recroom.com



Lewis Mbae

lewis@rudderstrack.com

rudderstack