

From 24h to 15min

How Robinhood Built a Streaming Lake House to Bring Data Freshness from 24h to <15min

Balaji Varadarajan

Senior Staff Engineer, Robinhood Markets

Apache Hudi PMC

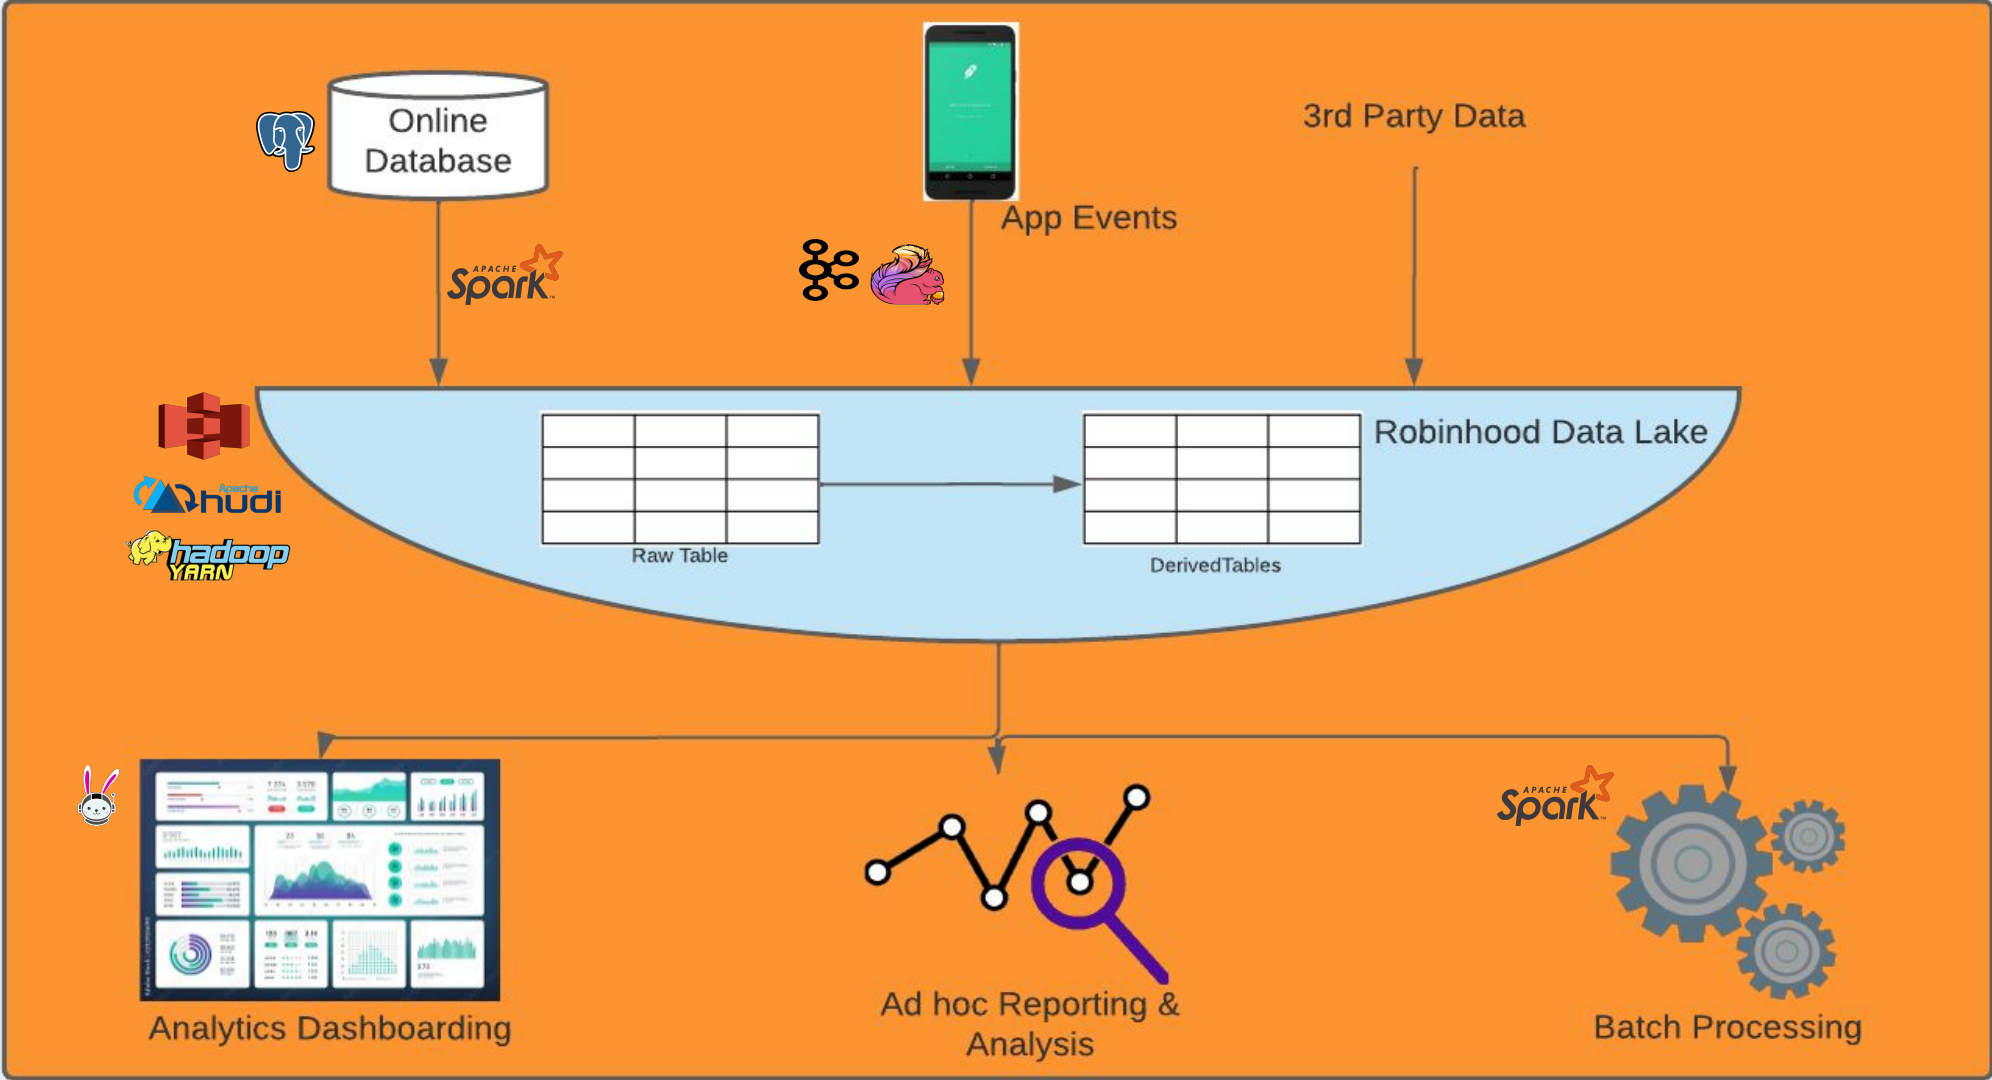
Vikrant Goel

Engineering Manager, Robinhood Markets

Agenda

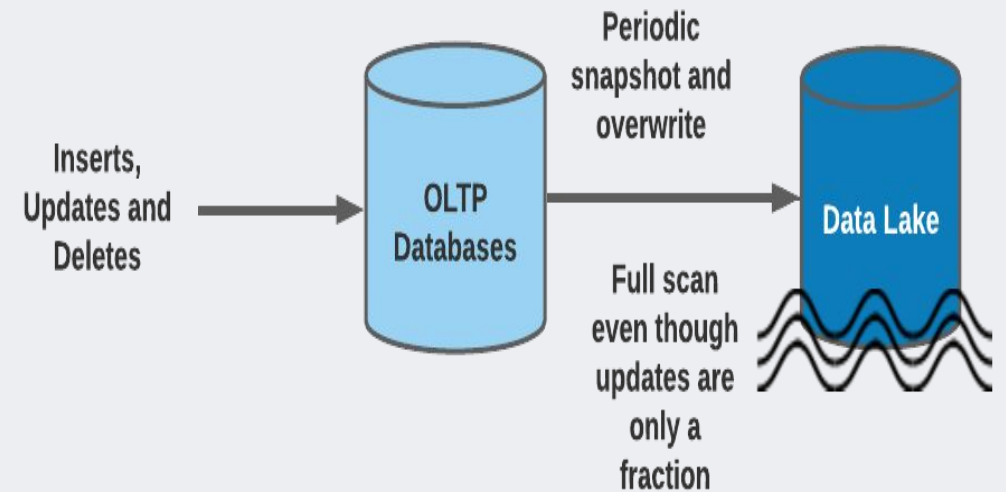
- Legacy Data Lake and Ingestion Framework
- Deep Dive - Change Data Capture (CDC)
 - Design
 - Lessons Learned
- Deep Dive - Data Lakehouse Ingestion
 - Apache Hudi
 - End to End Setup

Data Lake




Daily Snapshots

- Daily snapshotting of tables in RDBMS (RDS)
- High Read & Write amplifications
- Dedicated Replicas to isolate snapshot queries
- Bottlenecked by Replica I/O
- 24+ hours data latency

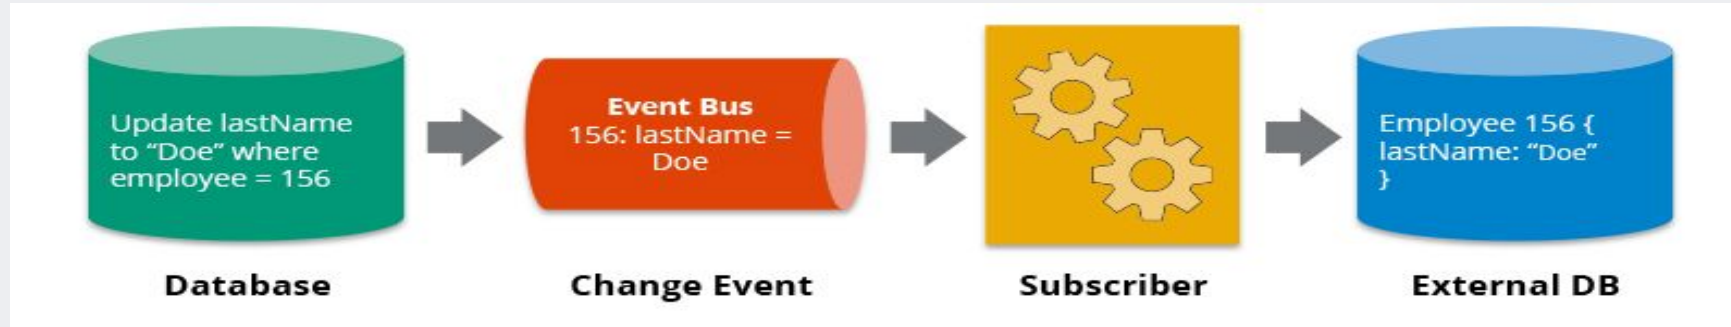


Need Faster Intraday Ingestion
Pipeline



*Unlock Data Lake for business critical
applications*

Change Data Capture



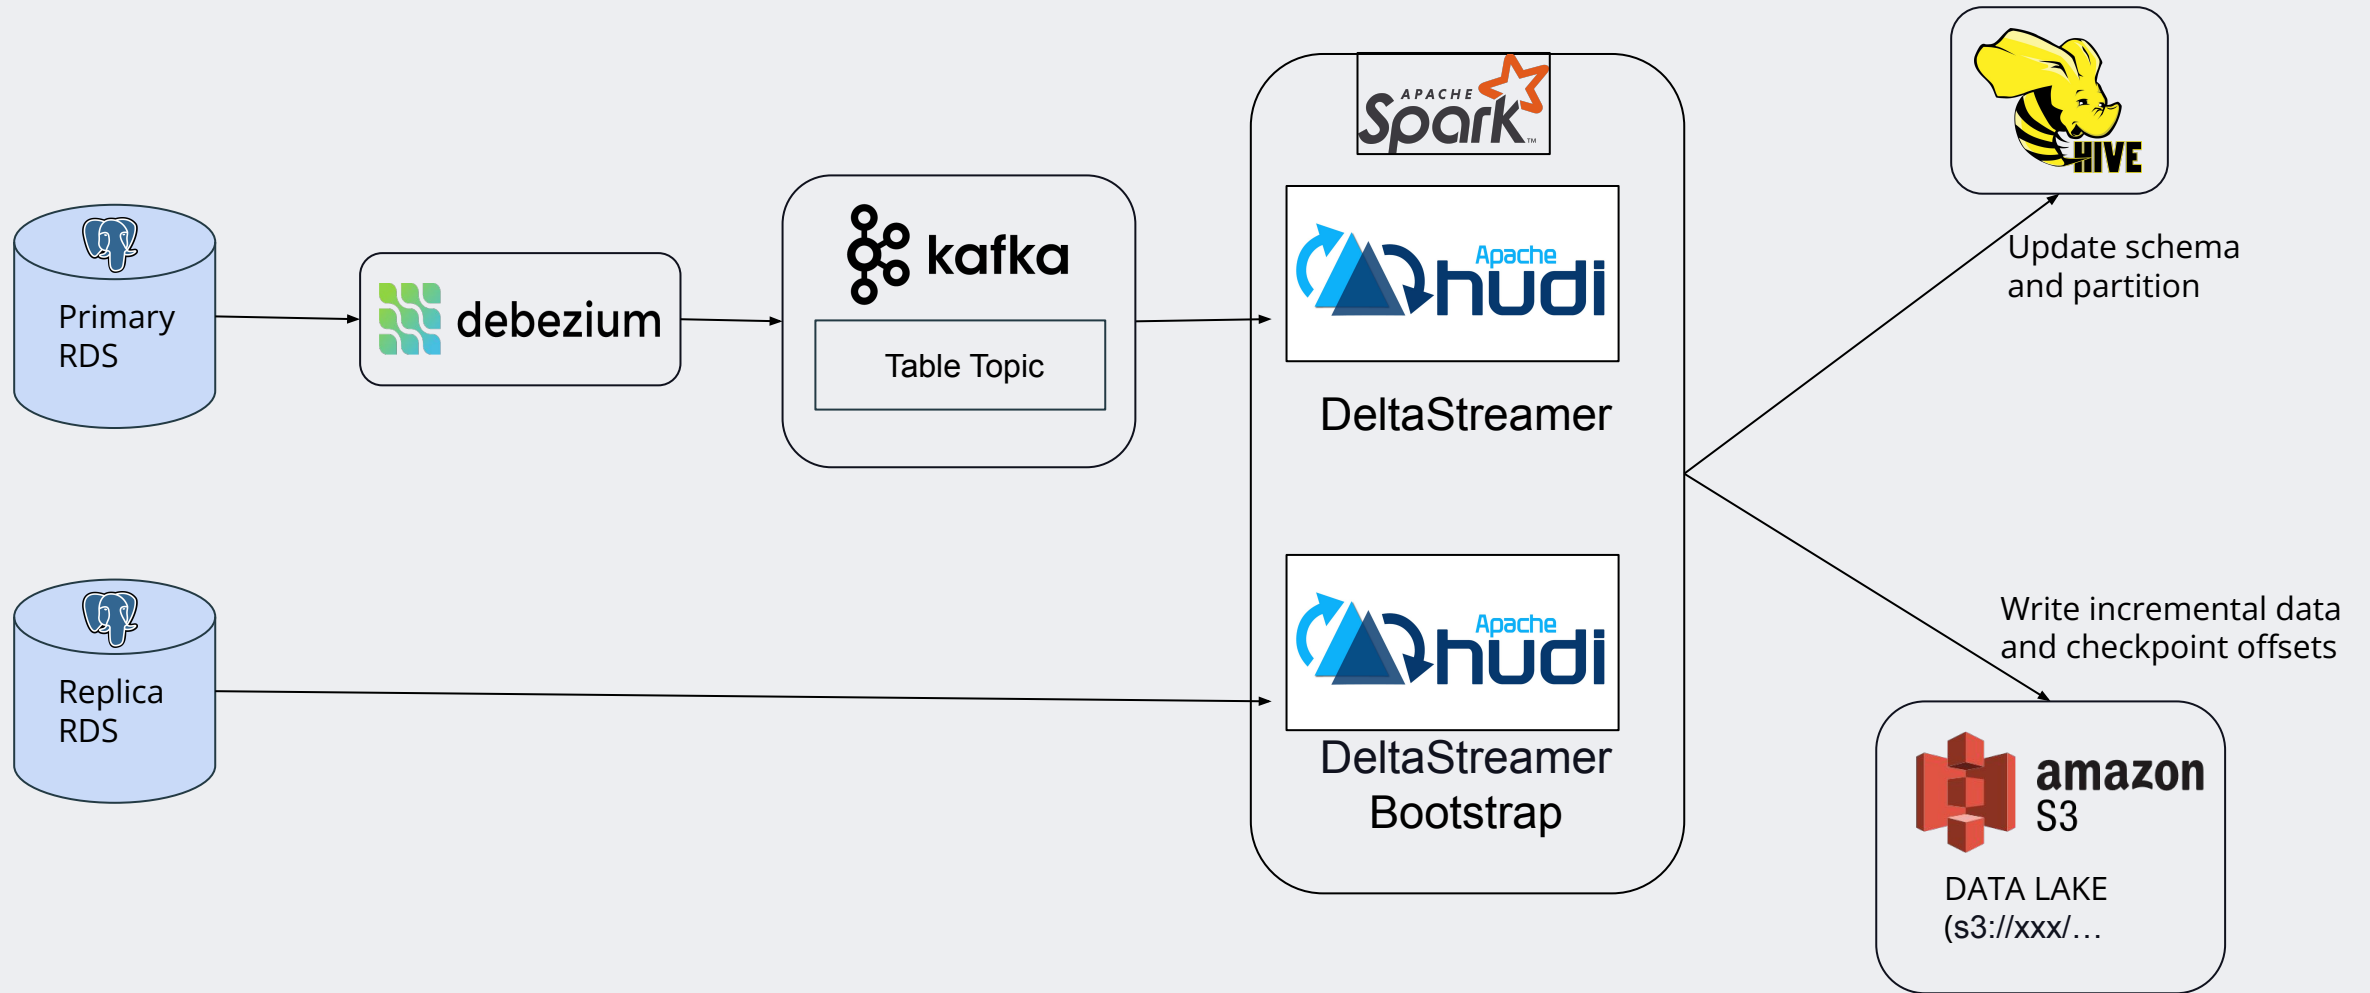
- Each CRUD operation streamed from DB to Subscriber
- Merge changes to lake house
- Efficient & Fast -> *Capture and Apply only deltas*

Deep Dive CDC using Debezium

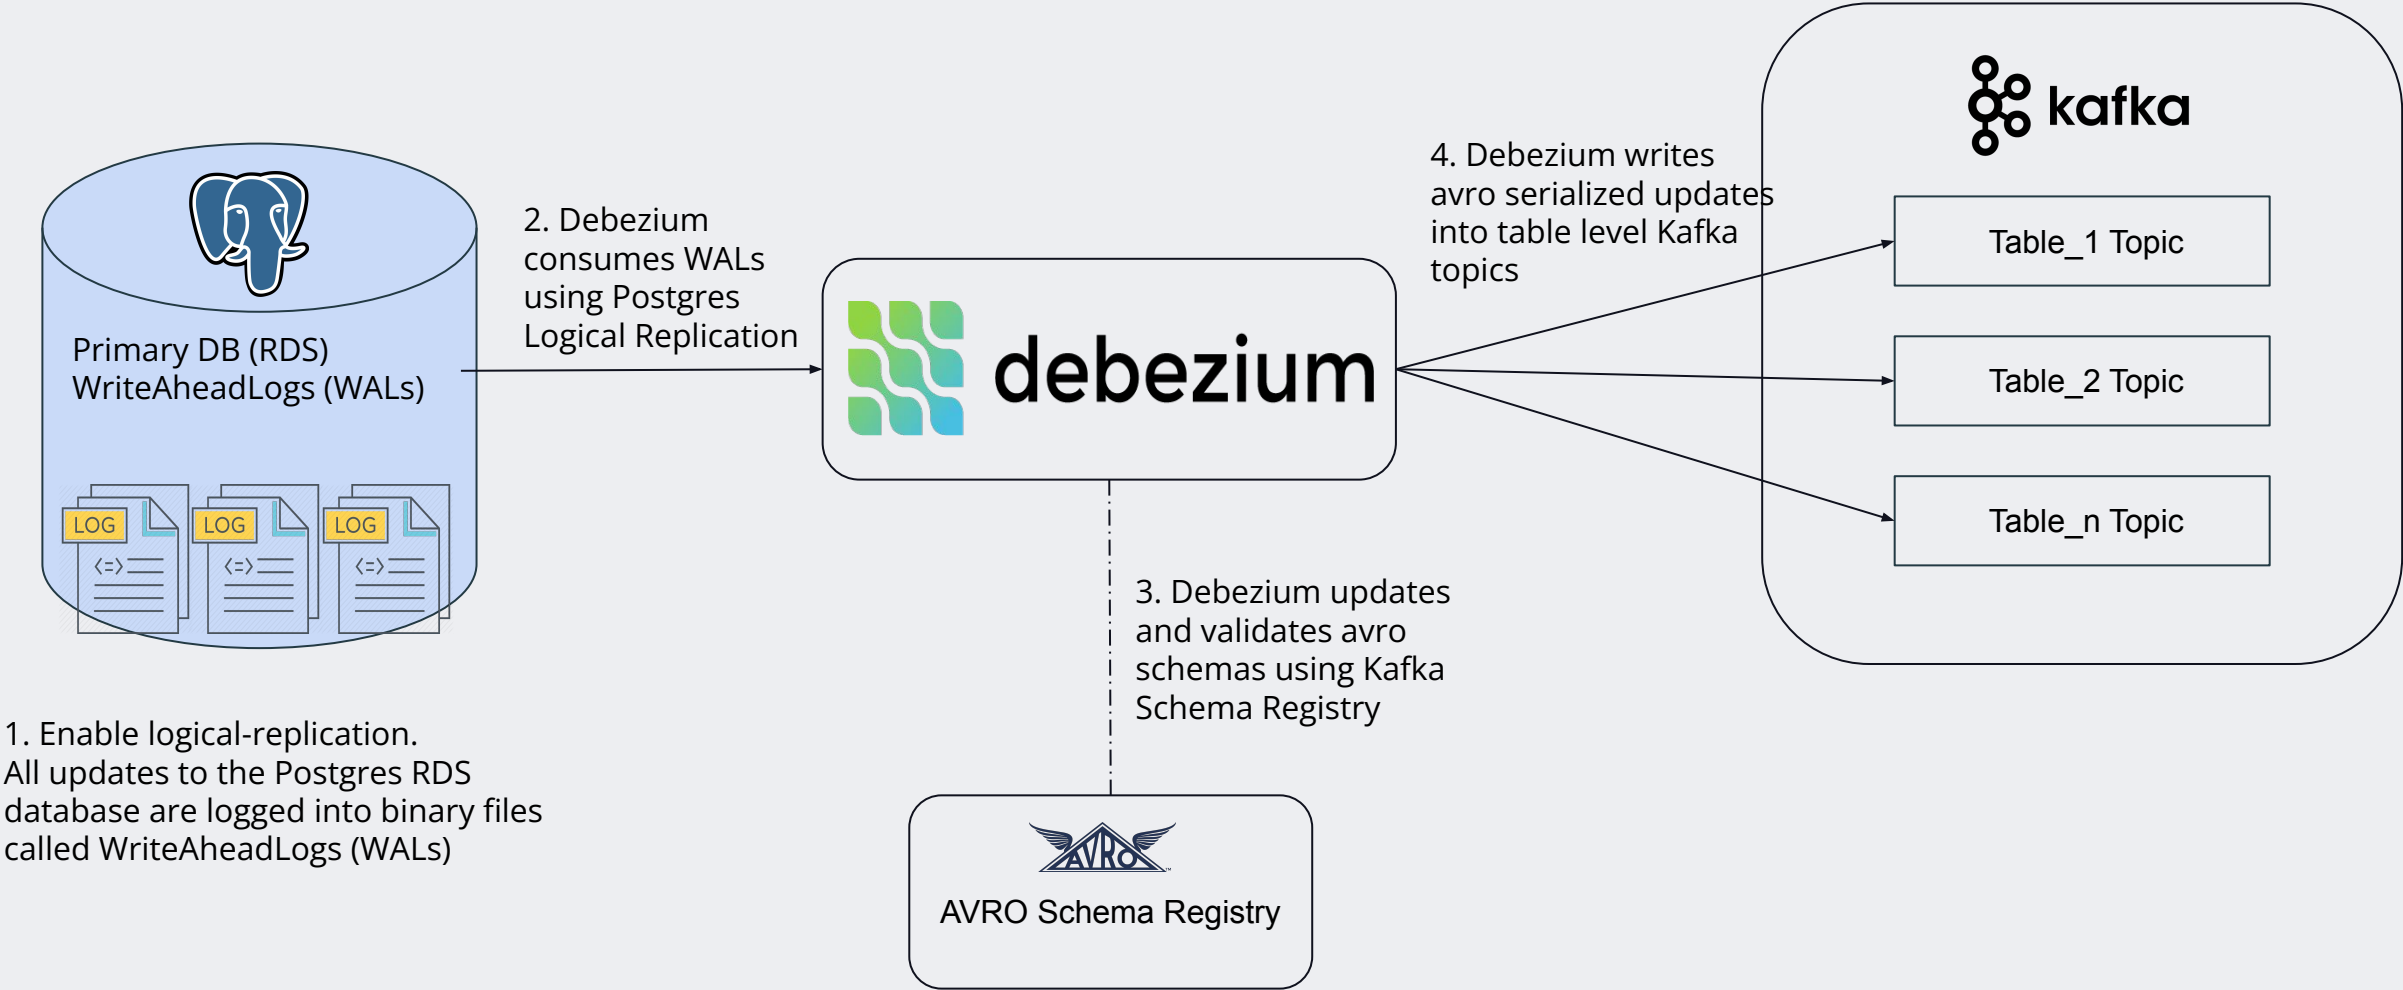
Debezium

- Open source & distributed Kafka-Connect Service for change data capture
- Support CDC from diverse RDBMS (Postgres, MySQL, MongoDB, etc.)
- Pluggable Sinks through Kafka

High Level Architecture



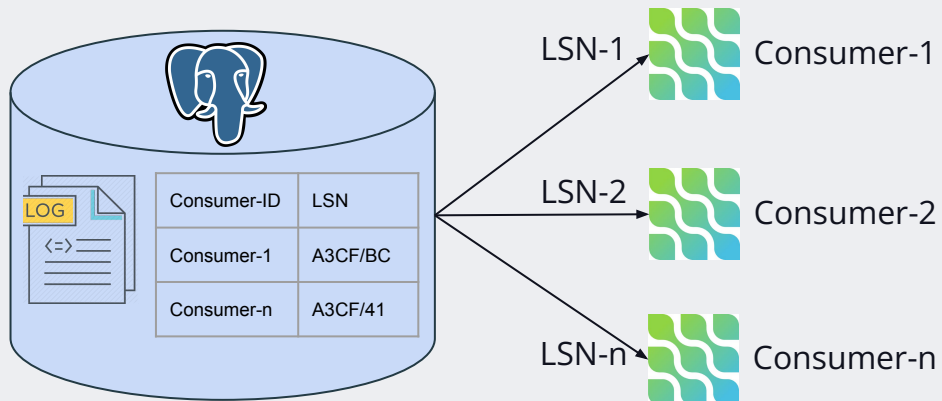
Debezium - Zooming In



Why did we choose Debezium over alternatives?

	Debezium	AWS Database Migration Service (DMS)
Operational Overhead	High	Low
Cost	Free, with engineering time cost	Relatively expensive, with negligible engineering time cost
Speed	High	Not enough
Customizations	Yes	No
Community Support	Debezium has a very active and helpful Gitter community.	Limited to AWS support.

Debezium: Lessons Learned



Postgres Primary:
- Publishes WALs
- Record LogSequenceNumber (LSN) for each consumer

1. Postgres Primary Dependency

ONLY the Postgres Primary publishes WriteAheadLogs (WALs).

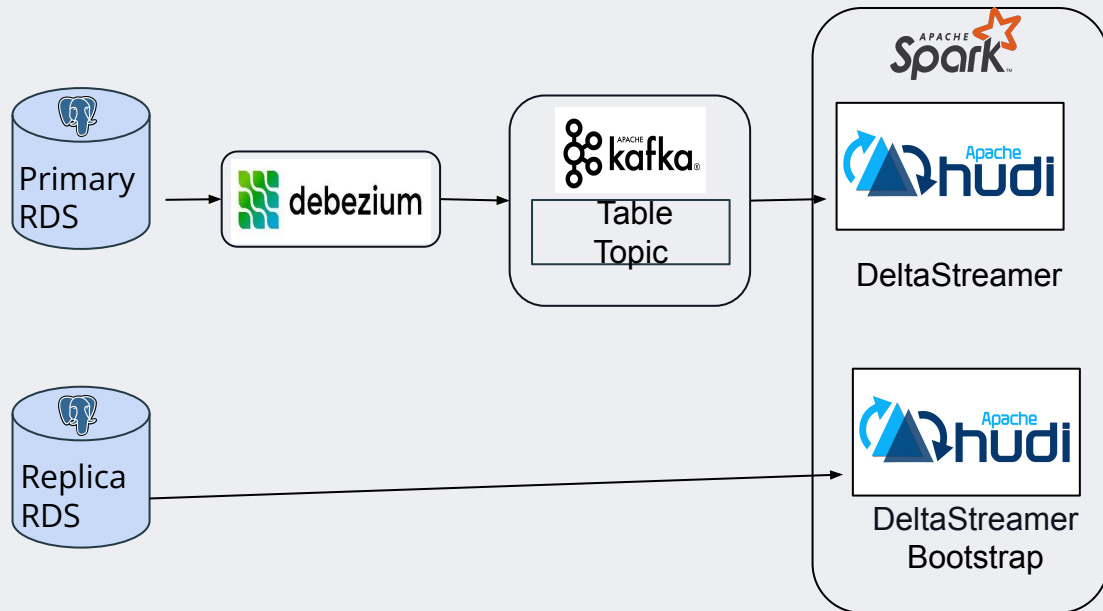
Disk Space:

- If a consumer dies, Postgres will keep accumulating WALs to ensure Data Consistency
- Can eat up all the disk space
- Need proper monitoring and alerting

CPU:

- Each logical replication consumer uses a small amount of CPU
- Postgres10+ uses pgoutput (built-in) : Lightweight
- Postgres9 uses wal2Json (3rd party) : Heavier
- Need upgrades to Postgres10+

Debezium: Lessons Learned



2. Initial Table Snapshot (Bootstrapping)

Need for bootstrapping:

- Each table to replicate requires initial snapshot, on top of which ongoing logical updates are applied

Problem with Debezium:

- Debezium processes data row at row level
- Large tables are **slow**
- Too much **pressure on Kafka Infrastructure and Postgres primary**

Solution using Hudi Deltastreamer:

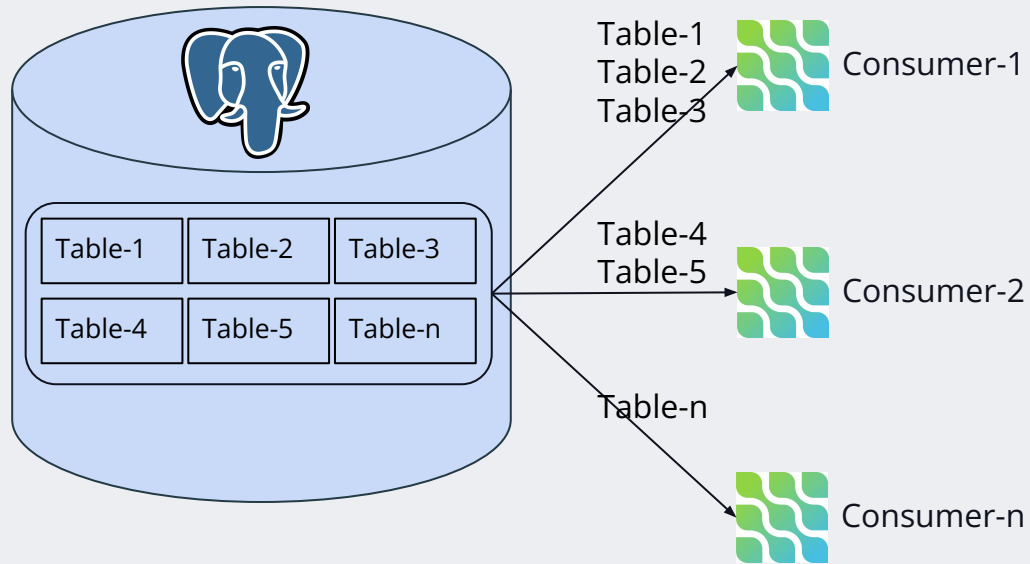
- Custom bootstrapping framework using **partitioned and distributed spark reads**
- **Can use read-replicas** instead of the master

Debezium: Lessons Learned

3. AVRO vs JSON

	AVRO	JSON	JSON + Schema
Throughput (Benchmarked using db.r5.24xlarge Postgres RDS instance)	Up to 40K mps	Up to 11K mps . JSON records are larger than AVRO.	Up to 3K mps . Schema adds considerable size to JSON records.
Data Types	<ul style="list-style-type: none">- Supports considerably high number of primitive and complex data types out of the box.- Great for type safety.	Values must be one of these 6 data types: <ul style="list-style-type: none">- String- Number- JSON object- Array- Boolean- Null	Same as JSON
Schema Registry	Required by clients to deserialize the data.	Optional	Optional

Debezium: Lessons Learned



4. Multiple logical replication streams for horizontal scaling

- Multiple large tables can overwhelm a single Debezium connector
- Split the tables across multiple Debezium connectors
Total throughput = throughput_per_connector * num_connectors
- Each connector does have small CPU cost

Debezium: Lessons Learned

5. Schema evolution and value of **Freezing Schemas**

Failed assumption: Schema changes are infrequent and always backwards compatible.

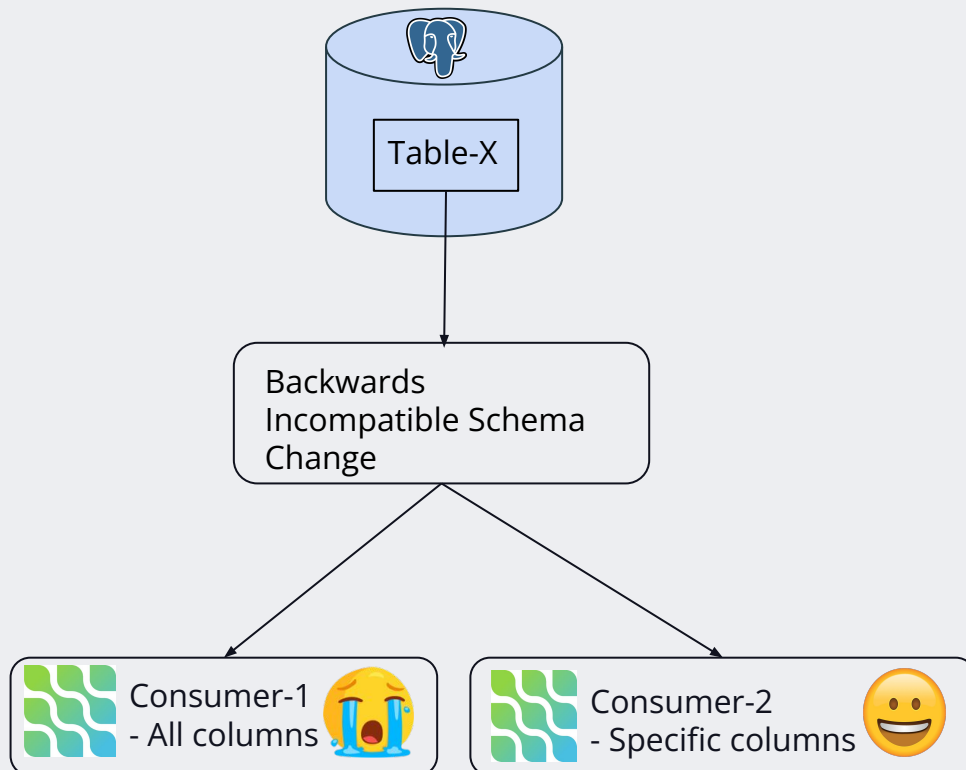
- Examples:
 1. Adding non-nullable columns (Most Common 99/100)
 2. Deleting columns
 3. Changing data types
 4. Can happen anytime during the day **#always_on_call**

How to handle the non backwards compatible changes?

- Re-bootstrap the table

Alternatives? Freeze the schema

- Debezium allows to specify the list of columns per table.
- Pros:
 - **#not_always_on_call**
 - Batch the changes for management window
- Cons:
 - Schema is temporarily out of sync



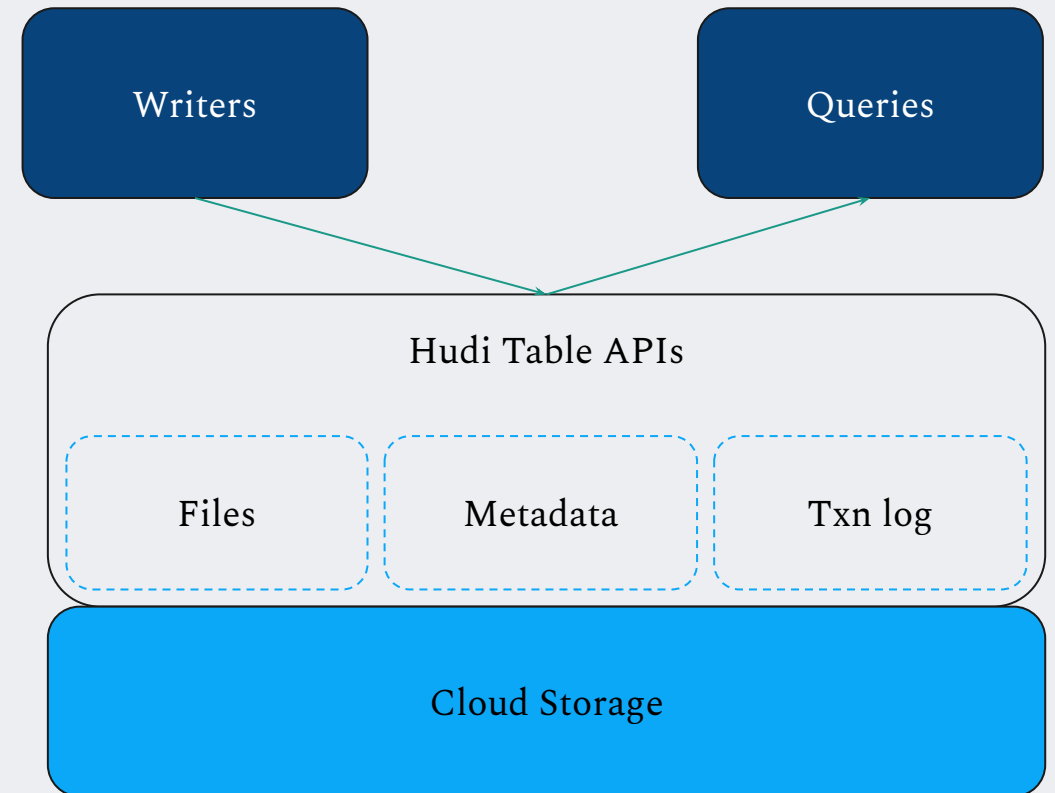
Deep Dive Lakehouse

Lakehouse - Requirements

- Transaction support
- Scalable Storage and compute
- Openness
- Direct access to files
- End-to-end streaming
- Diverse use cases

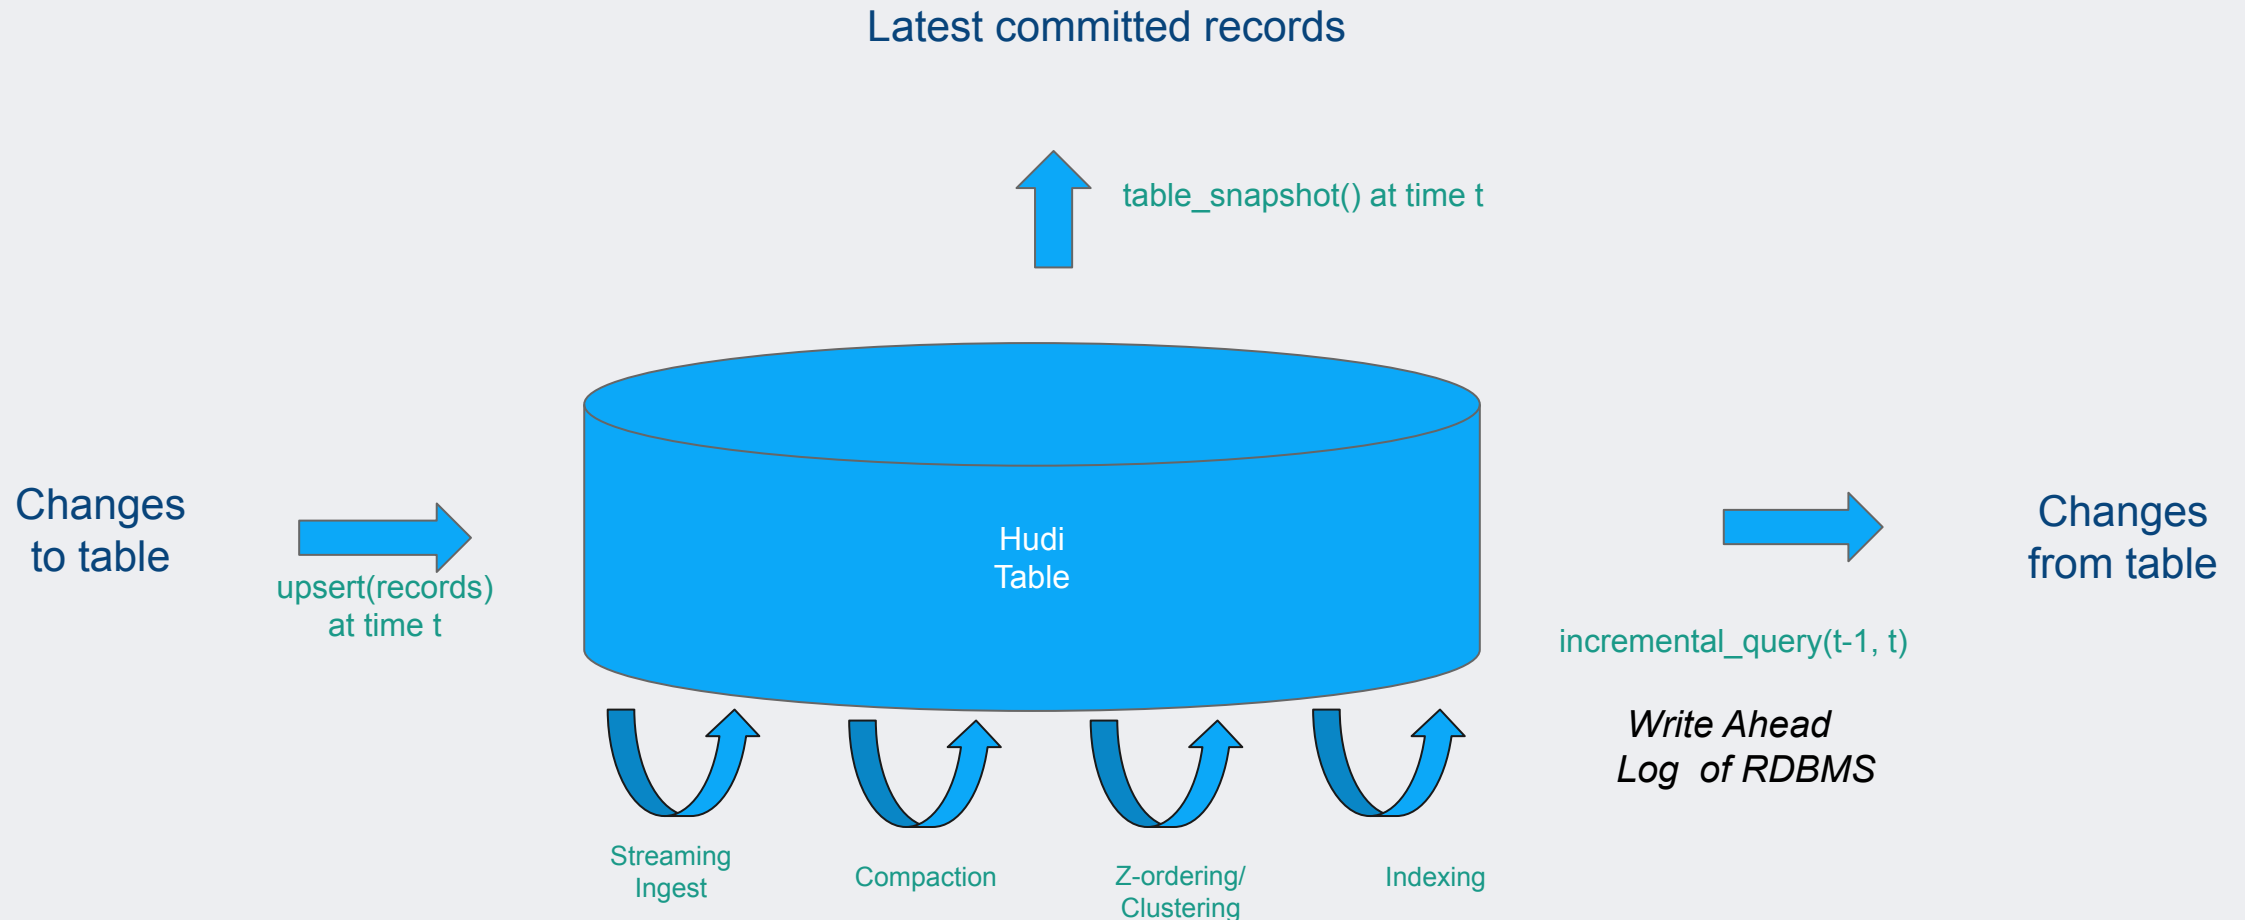
Apache Hudi - Introduction

- Transactional Lakehouse pioneered by Hudi
- Serverless, transactional layer over lakes.
- Multi-engine, Decoupled storage from engine/compute
- Upserts, Change capture on lakes
- Introduced Copy-On-Write and Merge-on-Read
- Ideas now heavily borrowed outside

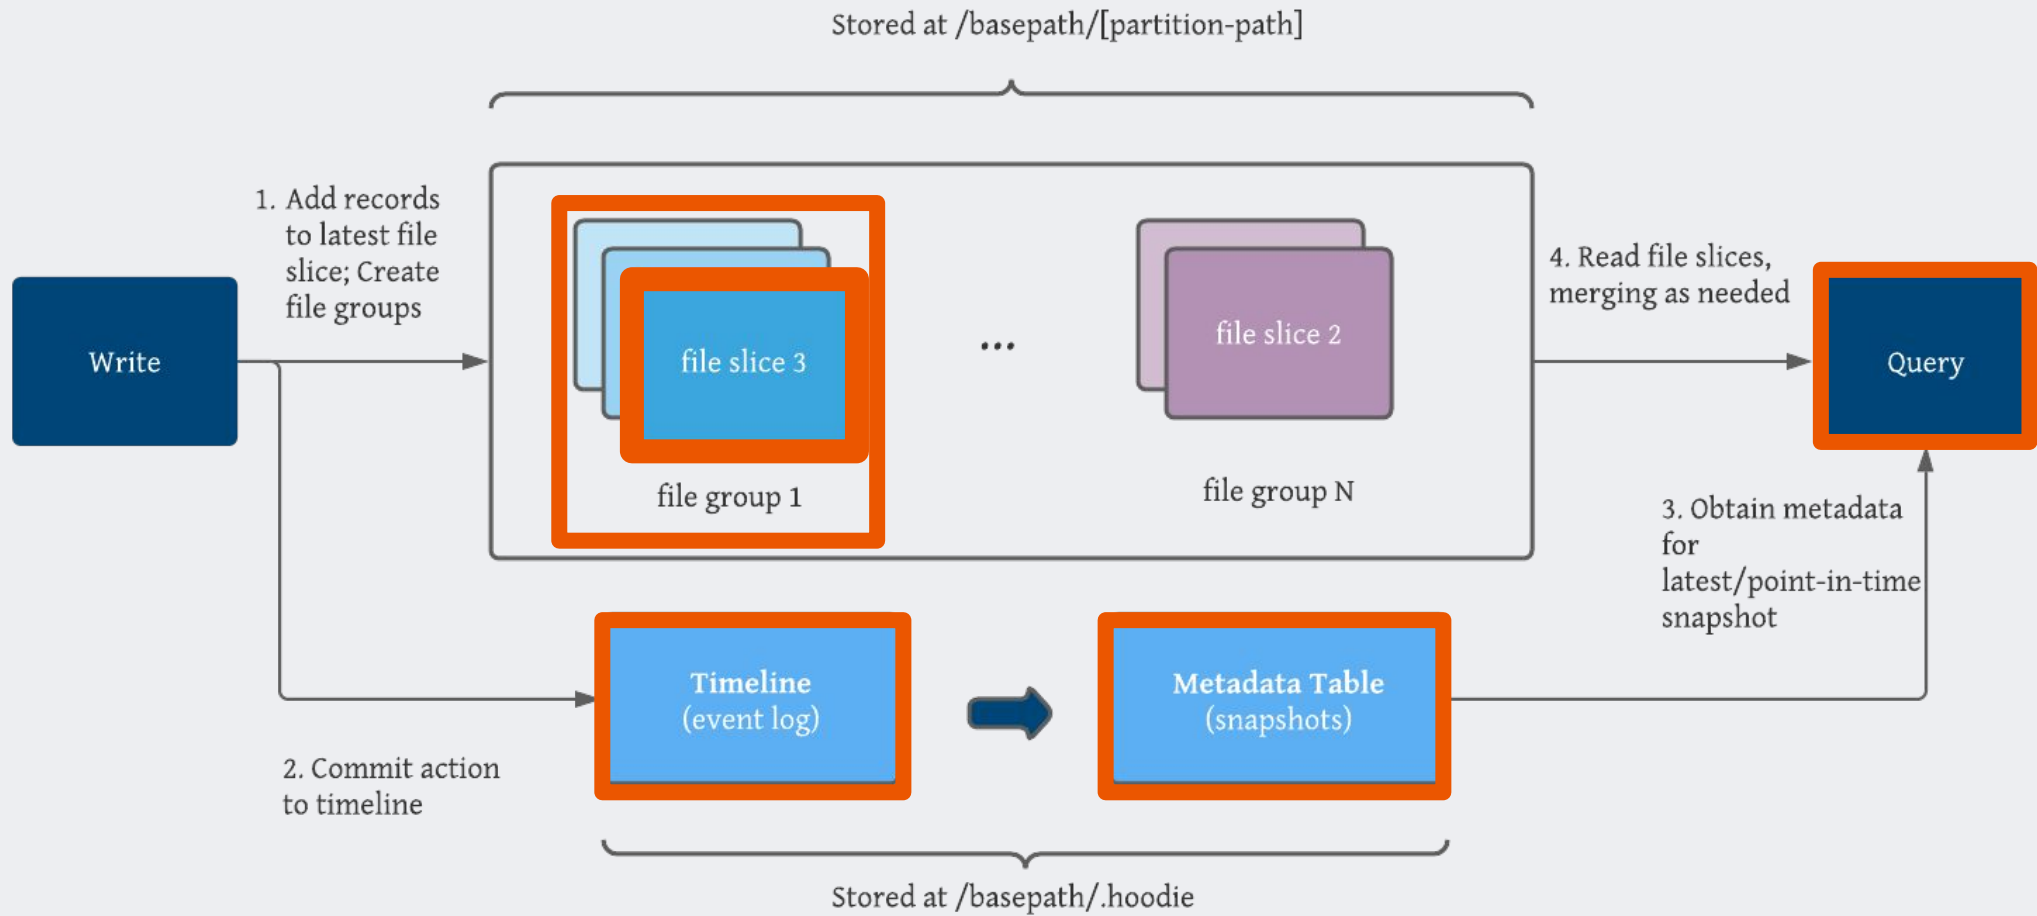


<https://eng.uber.com/hoodie/> Mar 2017

Apache Hudi - Upserts & Incrementals






Apache Hudi - Storage Layout



Apache Hudi - Copy-On-Write Table



Insert: A, B, C, D, E

commit time=0





Update: A => A', D => D'

commit time=1



Update: A' => A'', E => E', Insert: F

commit time=2



Snapshot Query

A,B,C,D,E

A',B,C,D',E

A'',B,C,D',E',F

Incremental Query

A,B,C,D,E


A',D'

A'',E',F

Apache Hudi - Merge-On-Read Table

Insert: A, B, C, D, E

commit time=0



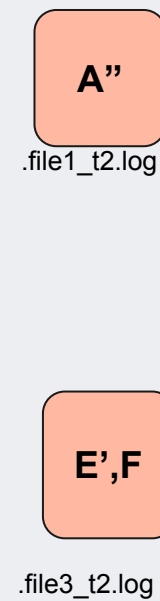
Update: A => A',
D => D'

commit time=1



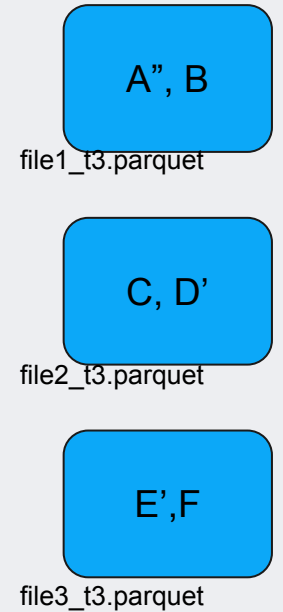
Update: A'=>A'',
E=>E', Insert: F

commit time=2



Compaction

commit time=3



Snapshot Query	A,B,C,D,E	A',B,C,D',E	A'',B,C,D',E',F	A'',B,C,D',E',F
Incremental Query	A,B,C,D,E	A',D'	A'',E',F	A'',E',F
Read Optimized Query	A,B,C,D,E	A,B,C,D,E	A,B,C,D,E	A'',B,C,D',E',F



The Community

Pre-installed on 5 cloud providers




Diverse PMC/Committers



Rich community of participants

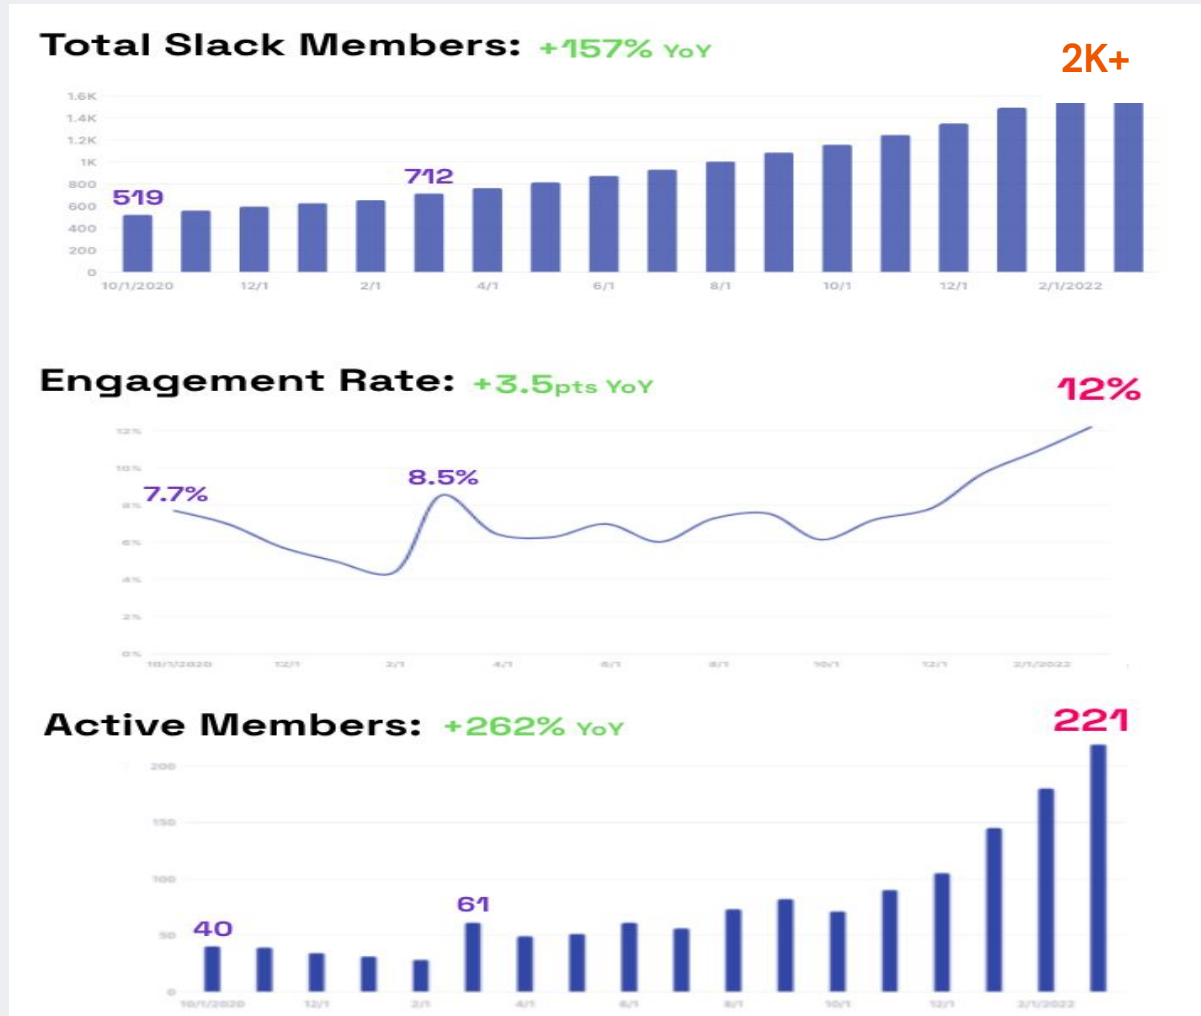


800B+ Records/Day <small>(from even just 1 customer!)</small>	2000+ Slack Members	1M DLs/month (400% YoY)
1000+ GH Engagers	225+ Contributors	20+ Committers



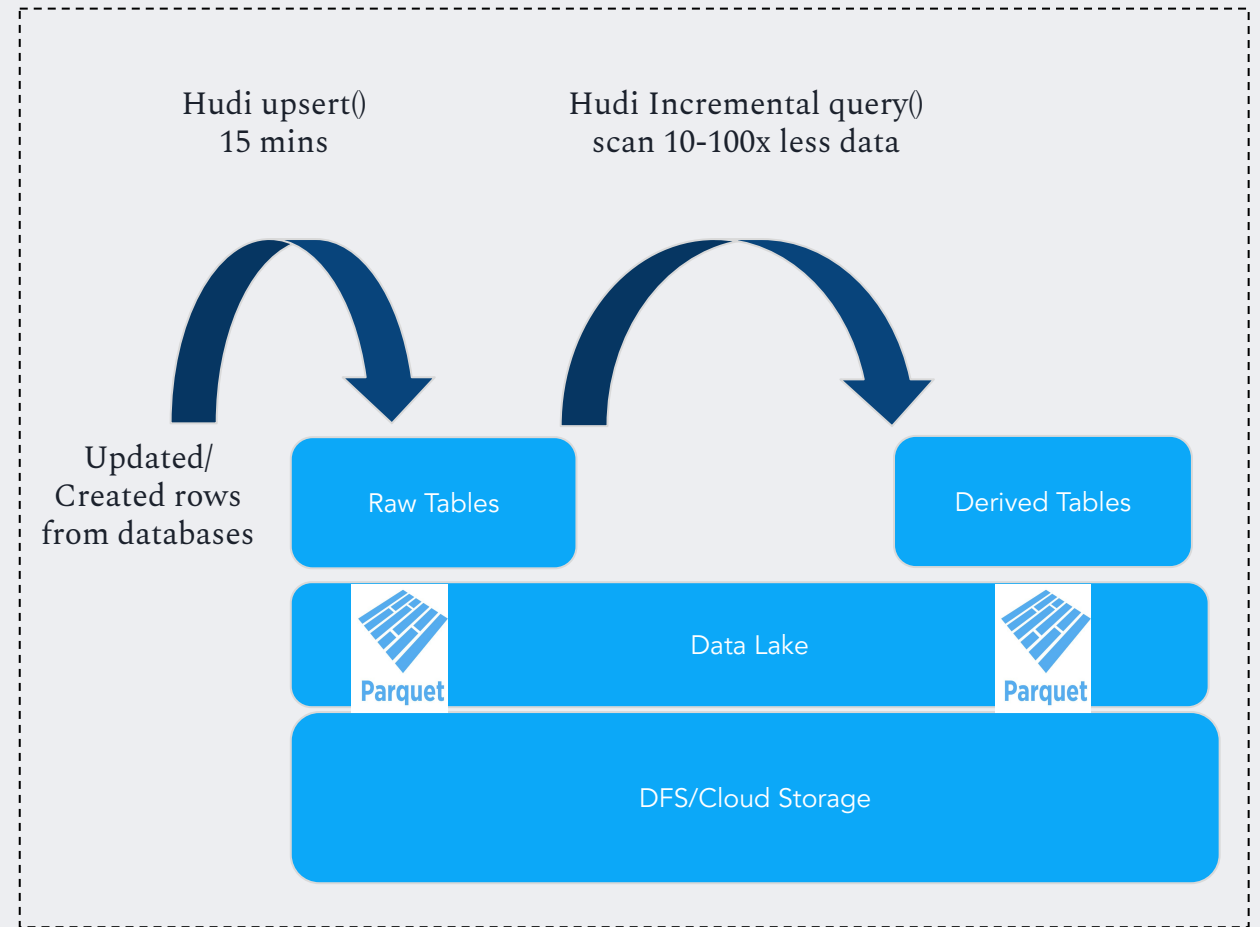
All trademarks, logos and brands are the property of their respective owners.

Apache Hudi - Community



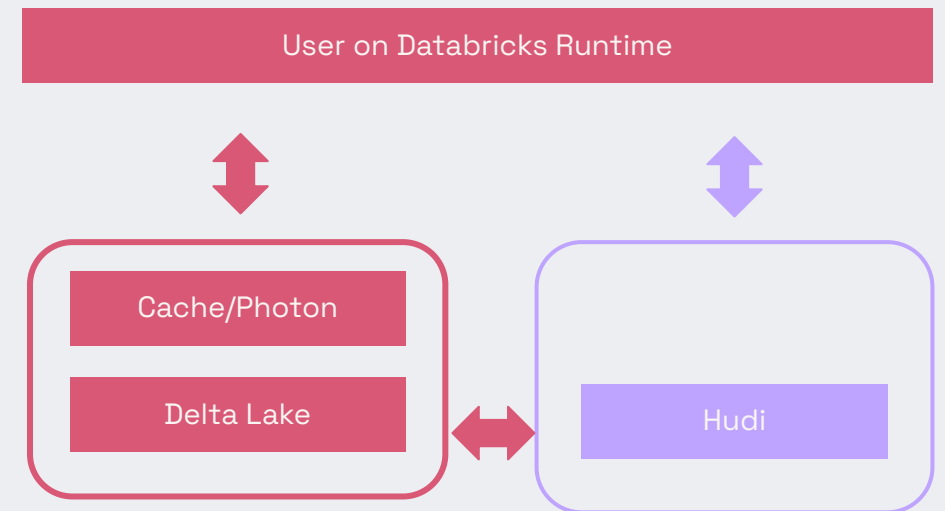
Apache Hudi - Relevant Features

- Database abstraction for cloud storage/hdfs
- Near real-time ingestion
- Incremental, Efficient ETL downstream
- ACID Guarantees



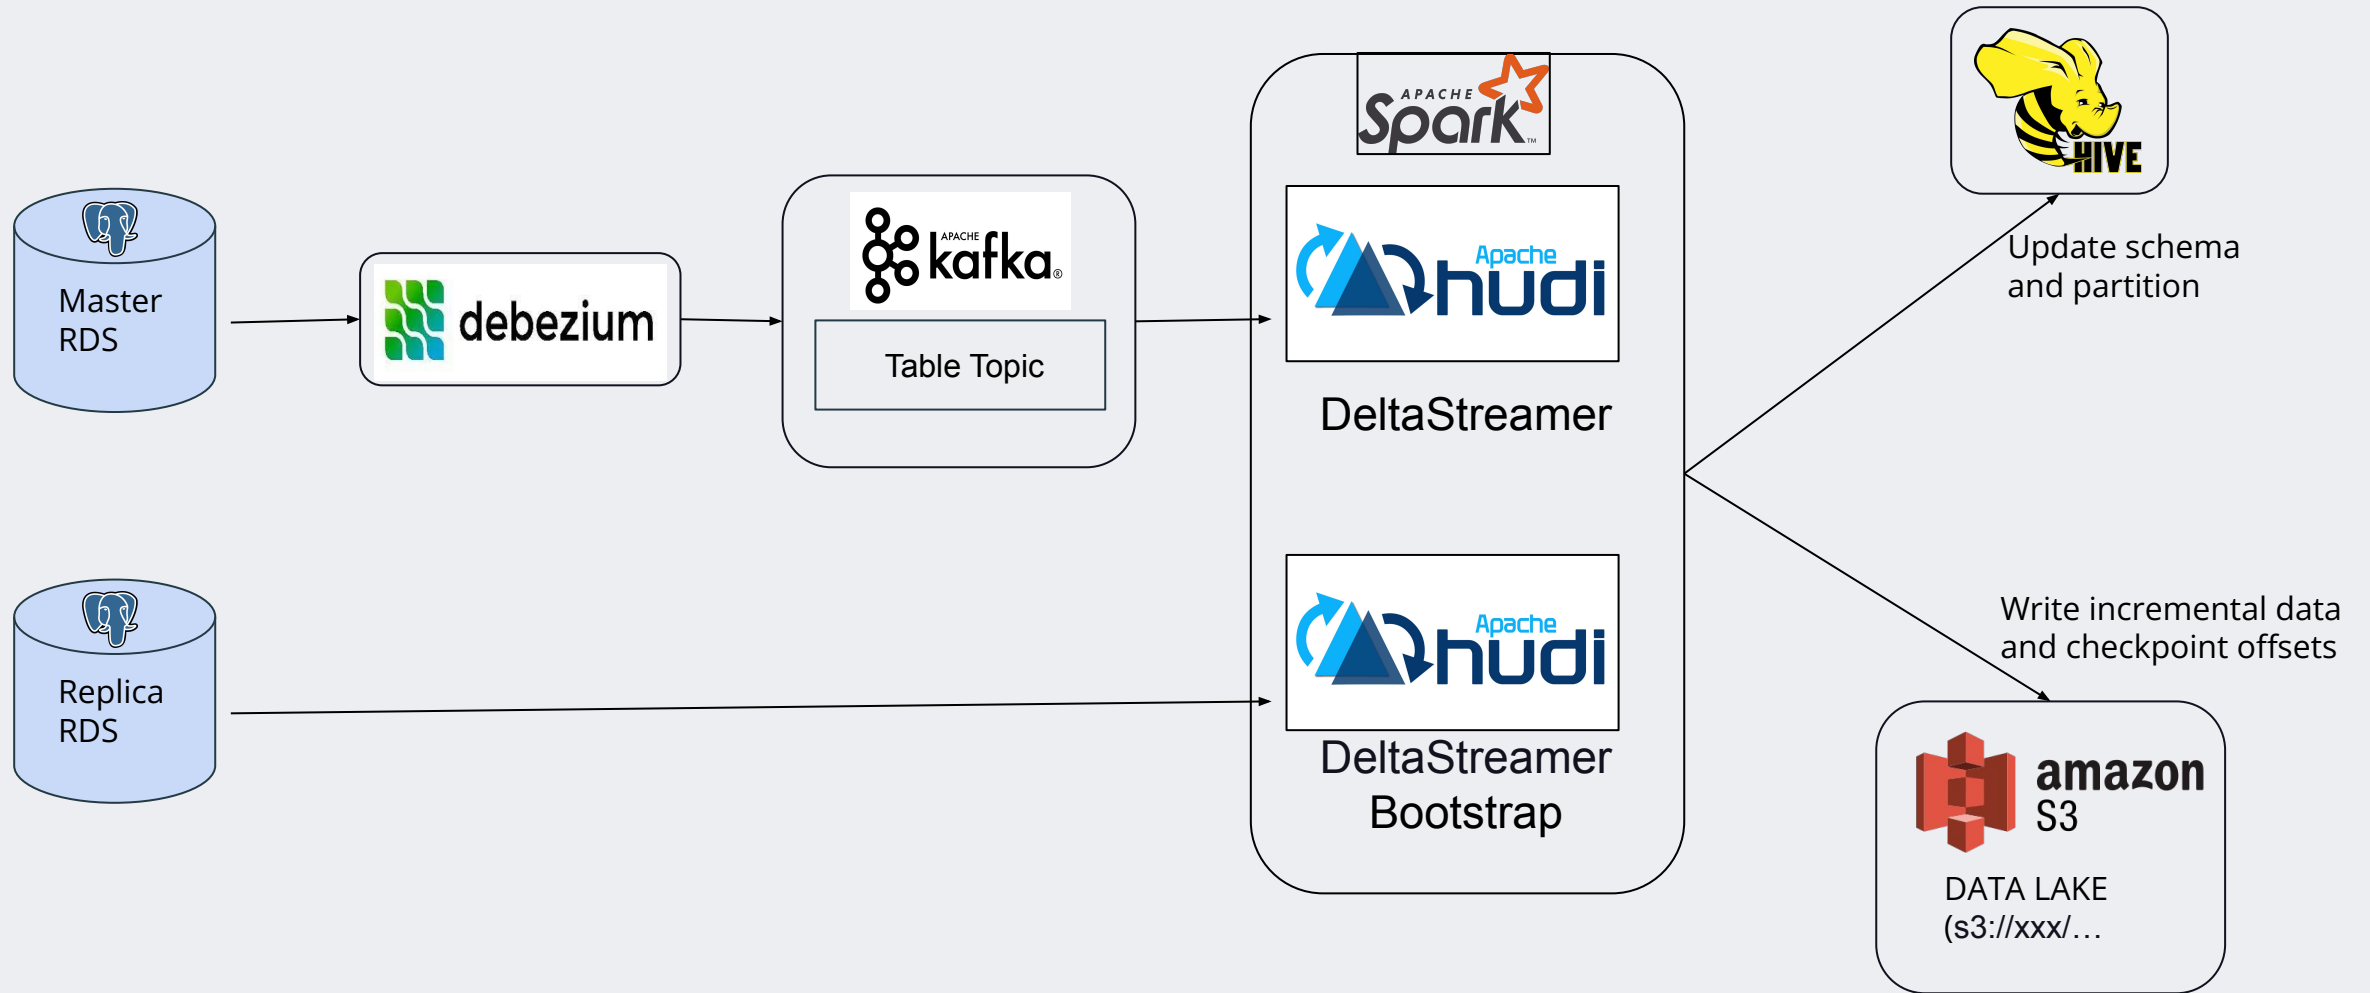
Apache Hudi + Databricks = Best of Both Worlds!

- We believe
 - Hudi offers the most complete lakehouse storage platform.
 - Databricks offers great Spark experience, but cache/photon don't work natively on Hudi.
- Hudi supports pluggable metadata syncing - BigQuery, Hive, Glue, DataHub, ...
- WIP - A bridge between Hudi and Delta Lake.
 - Query Hudi table snapshot as a Delta Lake table

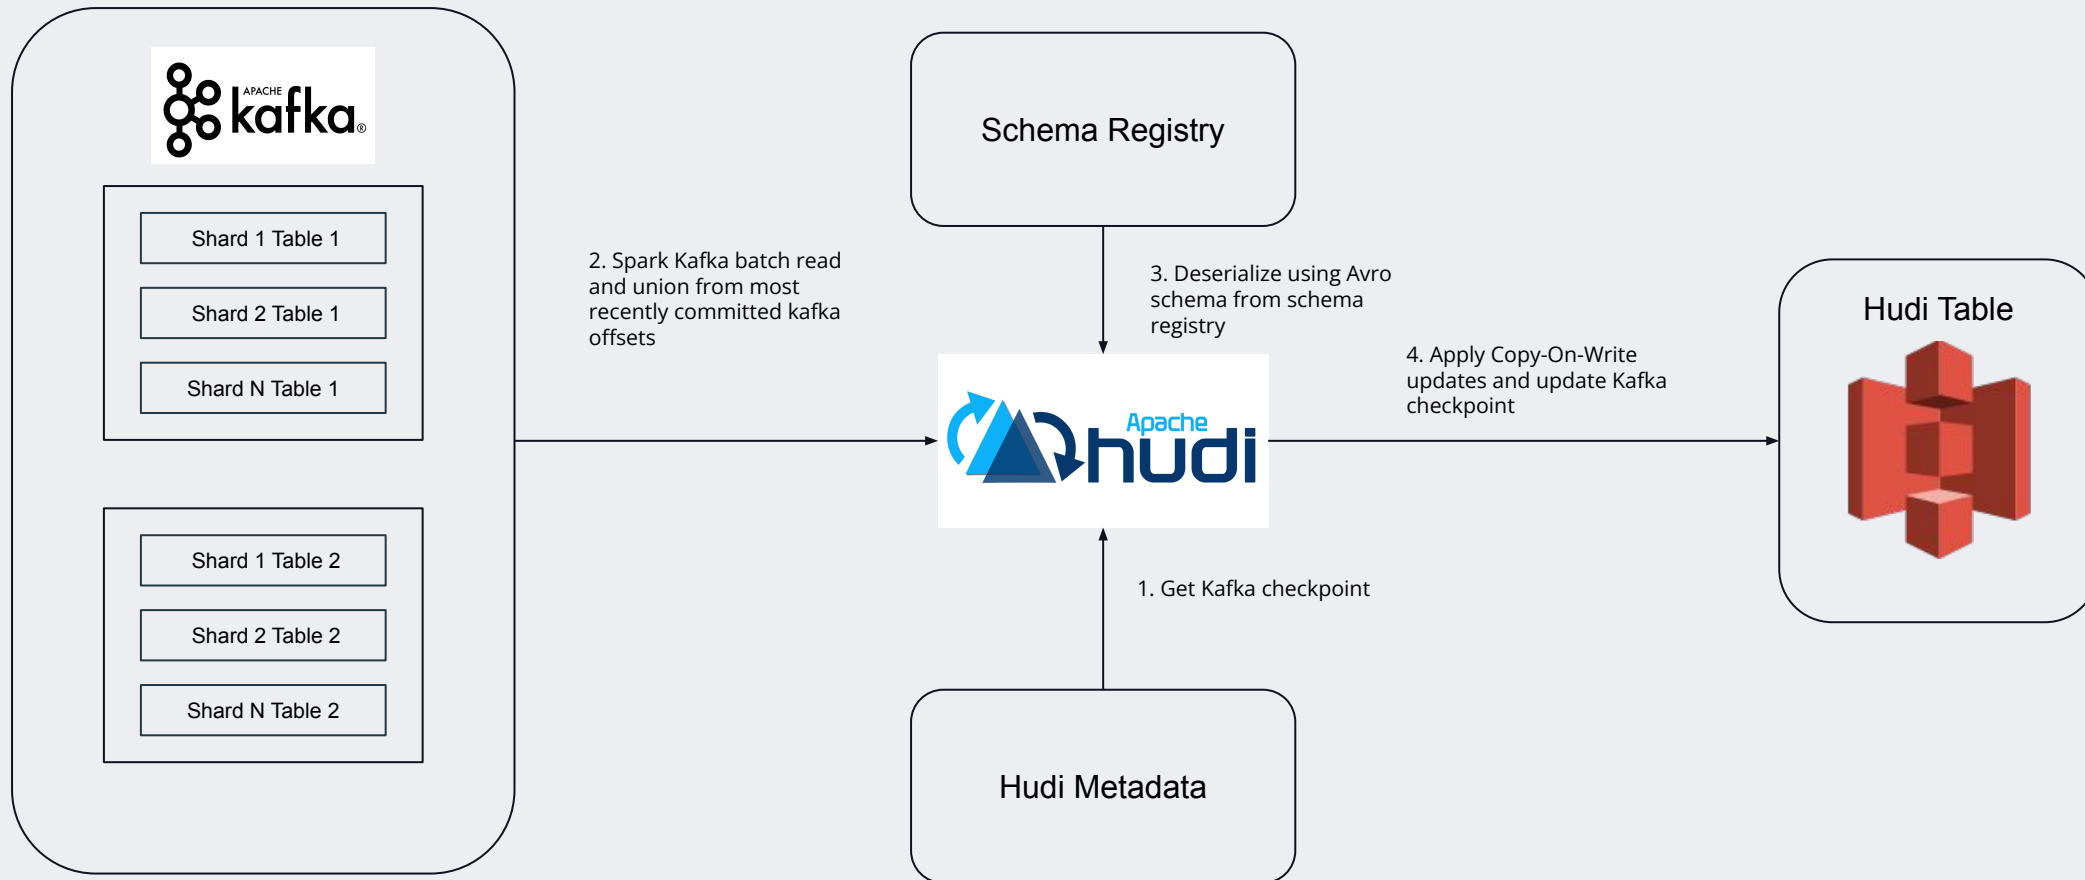


Deep Dive Ingestion

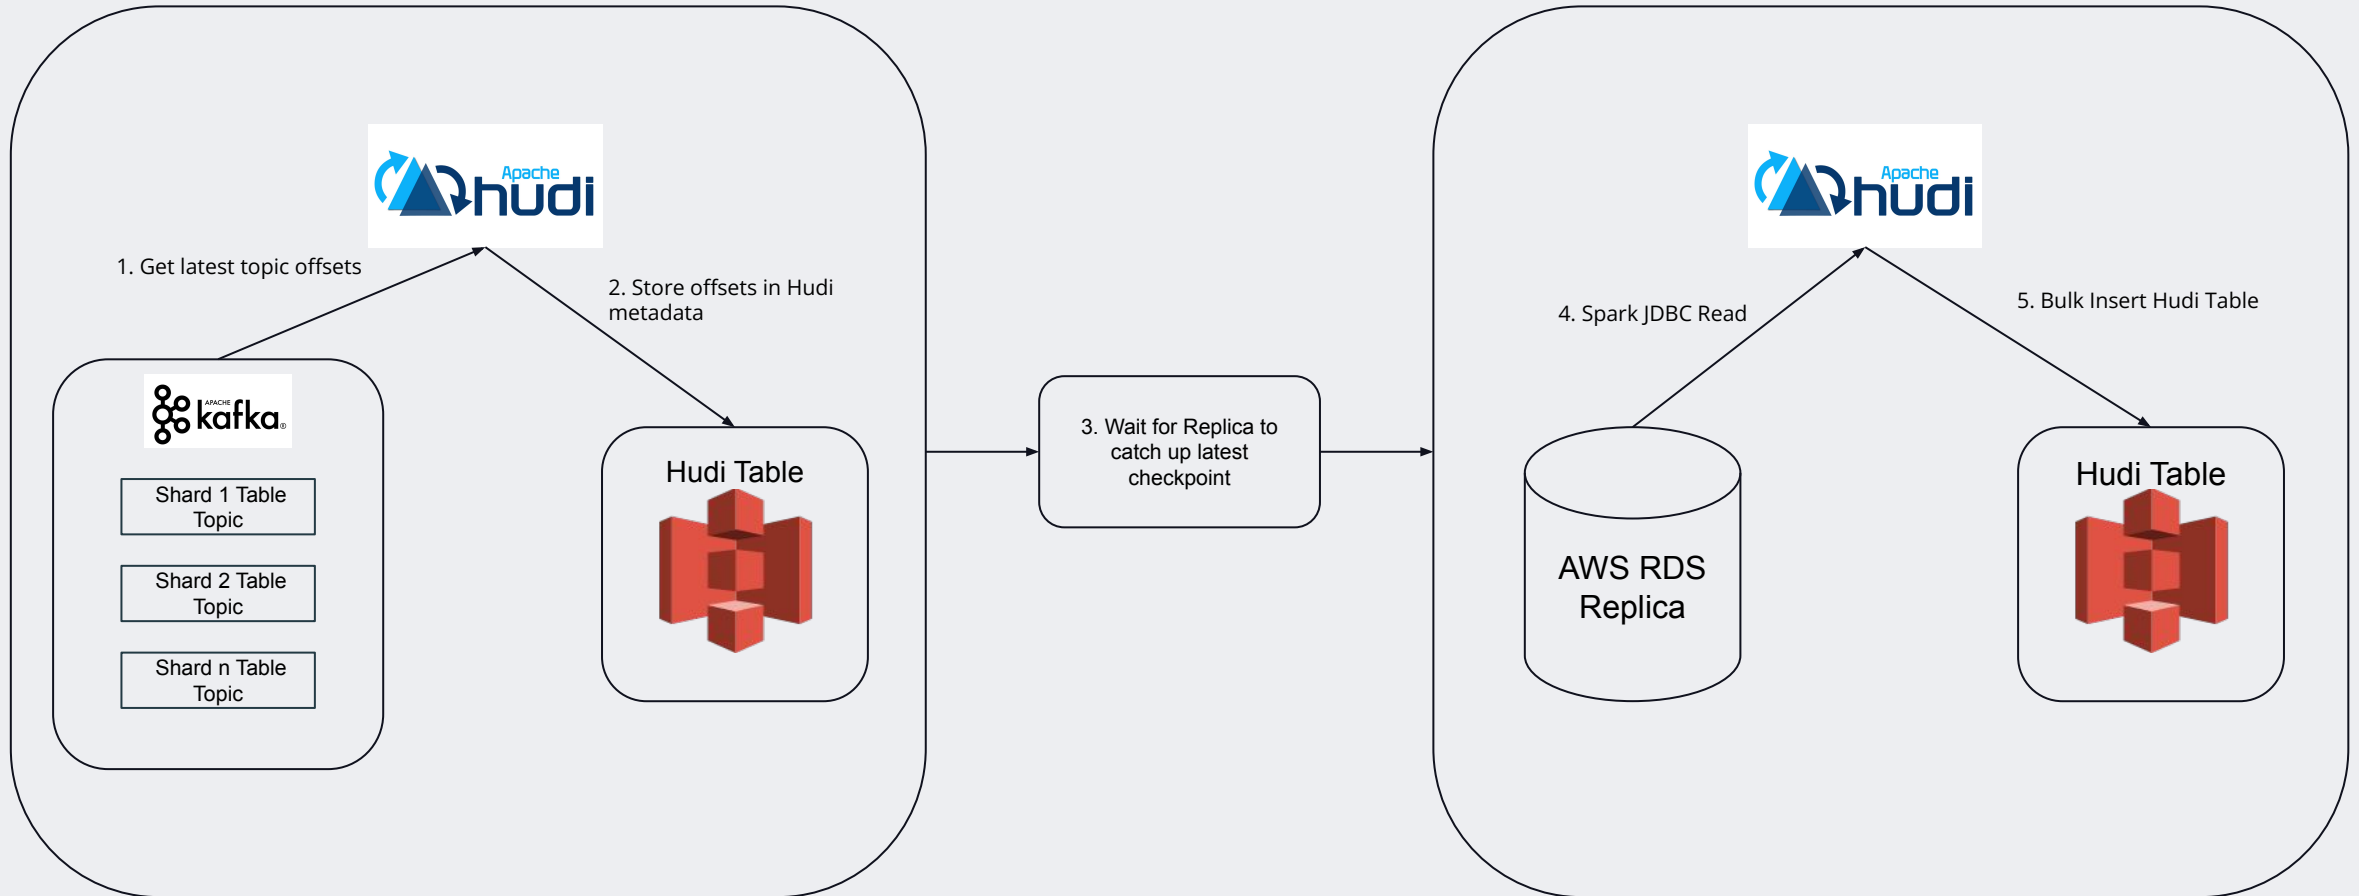
Recap - High Level Architecture



Data Lake Ingestion - CDC Path




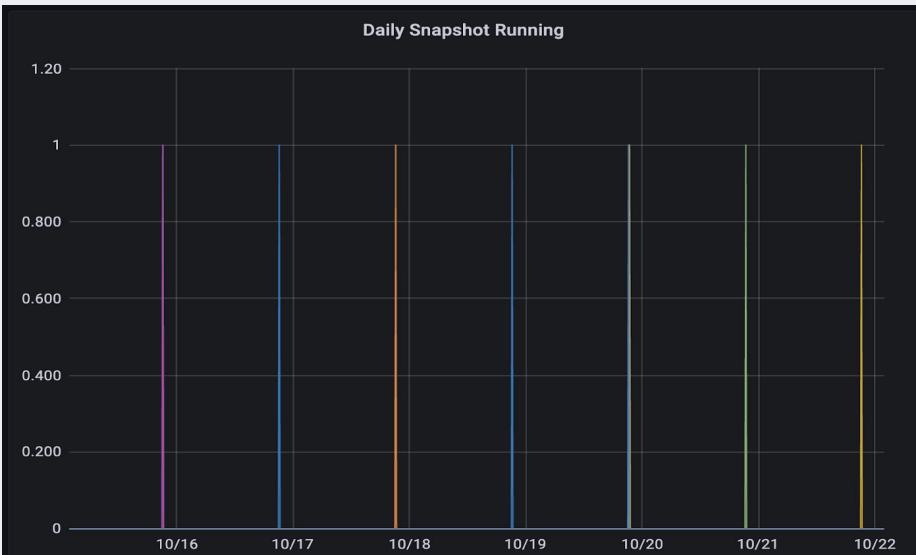
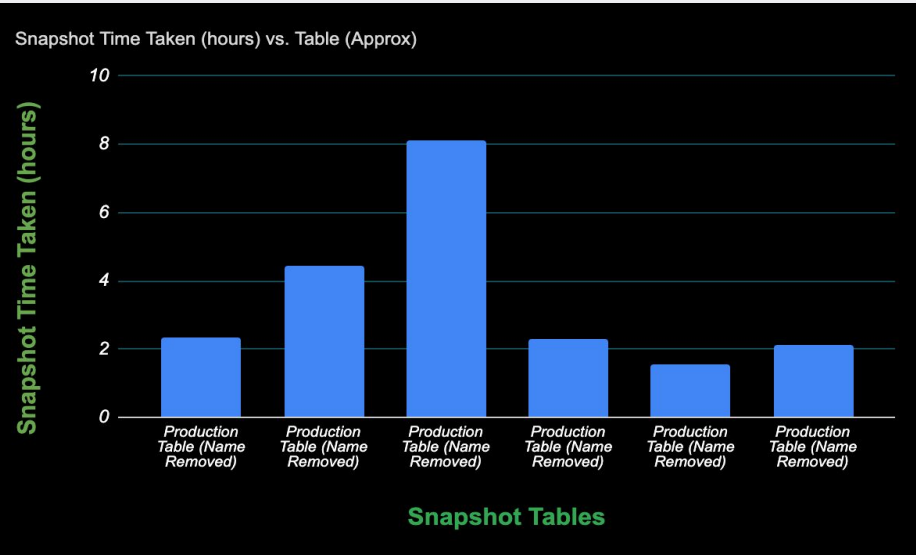
Data Lake Ingestion - Bootstrap Path



Running Ingestion at Scale

- 4000+ tables
- Automation, Self Healing
- Tiered SLAs - Provisioning and Isolation
- Pre-Commits and Validation for Quality Checks
- Monitoring & Alerting

Improved Freshness



DATA+AI
SUMMIT 2022

Thank you!

We'll hangout at back of room for any QA.

Balaji Varadarajan

Senior Staff Engineer, Robinhood Markets
Apache Hudi PMC

Vikrant Goel

Engineering Manager, Robinhood Markets