



PG Change Data Capture with Debezium

Hannu Valtonen

- Aiven co-founder
- Maintainer of Open Source projects: PGHoard, pglookout and pgmemcache
- PostgreSQL user for last 18 years



CDC Definition

In databases, Change Data Capture (CDC) is a set of software design patterns used to determine (and track) the data that has changed so that action can be taken using the changed data.

- Wikipedia



CDC - The Why

- Data's journey through your company's systems usually just starts with its initial storing
- Real-time change information stream as it happens
- \bigcirc
- No need to do bulk updates anymore with all their assorted errors
- Much more efficient, way fewer resources required since only delta is transferred



CDC - Why Apache Kafka

- Apache Kafka is <u>meant</u> for streaming data
- Huge ecosystem of tools to handle data streams
- Reliable
- Scalable
- Natural "message bus" for data from different databases

Foreword on examples

```
CREATE TABLE source_table (
    id SERIAL PRIMARY KEY,
    important_data text NOT NULL,
    create_time TIMESTAMPTZ NOT NULL DEFAULT clock_timestamp(),
    update_time TIMESTAMPTZ NOT NULL DEFAULT clock_timestamp(),
    updated BOOLEAN NOT NULL DEFAULT FALSE
);
ALTER TABLE source_table REPLICA IDENTITY FULL;
INSERT INTO source_table (important_data)
    VALUES ('first bit of very important analytics data');
INSERT INTO source_table (important_data)
        VALUES ('second bit of very important analytics data');
```

CDC - in the age of the Dinosaur

<Imagine a dinosaurs roaming freely through an idyllic landscape>

We're now talking about prehistoric times that predate the early 2000's

CDC - In the age of the Dinosaur

- Nightly database dump of some or all tables often done with pg_dump
- ETL from multiple databases to a single system
- Batch based
- Maybe using some proprietary ETL
- PostgreSQL COPY command made this less onerous



CDC - In the age of the Dinosaur

- Timestamp / sequence / status column based approach
- Add a column updated_timestamp to your table which you then read afterwards to try to find changed rows
- Same thing by having an updated boolean column in your table
- Possible limitations for noticing deletes or updates in naive implementations
- Confluent's Kafka JDBC connector works like this



CDC - in the age of the Dinosaur

```
SELECT * FROM source_table
WHERE id >= 0
ORDER BY id ASC LIMIT 1;
```

SELECT * FROM source_table WHERE timestamp >= y ORDER BY timestamp ASC LIMIT 1;

SELECT * FROM source_table WHERE updated IS FALSE ORDER BY id ASC LIMIT 1;

CDC - in the age of the Dinosaur

```
SELECT * FROM source_table
WHERE updated IS FALSE
ORDER BY id LIMIT 1;
UPDATE source_table SET updated = 't'
```

```
WHERE id = (SELECT id FROM source_table
WHERE updated IS FALSE
ORDER BY id ASC LIMIT 1)
RETURNING *;
```

CDC - Trigger based approaches

You create change tables that contain the INSERTed UPDATEd or DELETEd rows

- Slony, PGQ (Londiste) and plenty of homespun solutions
- Allows all DML (INSERTs, UPDATEs, DELETEs) to be extracted
- Bad performance as in doubles all writes done to the database
- Doesn't handle DDL (ALTER TABLE etc) gracefully



CDC - Trigger based approaches continued

- CREATE TRIGGER store_changes AFTER UPDATE, INSERT, DELETE ON source_table FOR EACH ROW EXECUTE PROCEDURE store_change();
- And then the trigger just INSERTs the contents of the change to a change table with the information stating whether it was an INSERT, UPDATE or DELETE
- The change table contents are read and applied from start to finish in some other database



CDC - Advent of a new age

<Imagine something very modern>

PG's built-in logical decoding saga started with the release of 9.4 at the end of '14

CDC - Logical decoding - What is it?

- PostgreSQL can keep track of all the changes happening in a database
- Decodes WAL to desired output format
- Multiple logical decoding output plugins exist
- Very performant, low-overhead solution for CDC
- Avoids the multiple write problem with triggers by using the WAL that PG was going to write anyway



CDC - Logical decoding - What can it do?

- Track all DML (INSERT, UPDATE, DELETE) changes
- Unit of Change is a row of data that's already committed
- Allows reading only the wanted subset of changes
- Use cases include auditing, CDC, replication and many more
- Logical replication connections supported in multiple PostgreSQL drivers (JDBC, Python psycopg2)



CDC - Logical decoding - What can't it do?

- Replicate DDL
- Possible to set up event triggers that write to a table and then have your replication system run the DDL based on it
- Depending on output plugin some data types not supported
- Failovers not handled gracefully as replication slots exist only on master nodes
- Changes tracked are limited to a single logical DB



CDC - Logical decoding - How to set it up?

```
postgresql.conf:
wal_level=logical
max_replication_slots = 10 # At least one
max_wal_sender = 10 # At least one
```

\$ CREATE ROLE foo REPLICATION LOGIN;

Before PG 10 also needs changes to pg_hba.conf

CDC - wal2json

- Author: Euler Taveira de Oliveira
- Decodes logical changes into JSON format
- Datatype limitations (JSON doesn't natively support everything)
- Supported by multiple DBaaS vendors (Aiven, AWS RDS, <u>https://github.com/eulerto/wal2json</u>)
- Supported by Debezium 0.7.x+



Debezium

- Apache Kafka Connect Connector plugin (http://debezium.io/)
- Uses logical replication to replicate a stream of changes to a Kafka topic
- Supports PostgreSQL, MongoDB, MySQL, (Oracle)
- Uses log compaction, only needs to keep the latest value (if you pre-create topics)
- Can run arbitrary transformation code on the data as it's received
- Supports protobuf output plugin or wal2json



Why Debezium

- Gets the data in real-time from PostgreSQL No more waiting
- Once you get the data to Kafka you can process it whichever way
- Plenty of other Kafka Connect connectors to send it to the next system
- Basis for Kafka centric architectures
- You don't need to know beforehand who is going to consume the data or why

Debezium gotchas

- Remember to set REPLICA IDENTITY FULL to see UPDATE, DELETE changes
- When PG master failover occurs, PG replication slot disappears
 - Need to recreate state
- If you don't pre-create topics they use DELETE not COMPACT as cleanup policy
- Limited datatype support
- Unlike documentation says, sslmode param is "require", not "required"



Debezium example

```
curl -H "Content-type:application/json" -X POST
https://avnadmin:zqv9z695oo5e1k3h@debezium-pg-demoproject.aivencloud.com:25649/connectors -d
' {
 "name": "test_connector",
 "config": {
   "connector.class": "io.debezium.connector.postgresgl.PostgresConnector",
   "database.hostname": "debezium-pg-demoproject.aivencloud.com",
   "database.port": "22737",
   "database.user": "avnadmin",
   "database.password": "ngj26a2lni8pi2ax",
   "database.dbname": "defaultdb",
   "database.server.name": "debezium-pg",
   "table.whitelist": "public.source_table",
   "plugin.name": "wal2json",
   "database.sslmode": "require"
}'
```

Demo

If we have time...



CDC - Recap

- Logical decoding and replication have revolutionized the way CDC can be done with PostgreSQL
- We're only seeing the very beginnings of its adoption
- Note that logical decoding is not a perfect solution (yet)
- Apache Kafka a natural fit it is <u>meant</u> for streaming data



Q & A

Time to ask me anything







The end

https://aiven.io

🔰 @aiven_io

aiven