An AWS DMS Replication Journey

… from Oracle to MySQL

MARIS ELSINS
Lead Database Consultant
Pythian

ORACLE ACE

Apress AUTHOR

ORACLE Certified Master

aws CERTIFIED Solutions Architect Associate

Google Cloud Certified PROFESSIONAL CLOUD ARCHITECT

Maris Elsins

Lead Database Consultant

At Pythian since 2011

Located in Riga, Latvia

Oracle [Apps] DBA since 2005

Speaker at conferences since 2007

@MarisDBA

elsins@pythian.com

ABOUT PYTHIAN

Pythian’s 400+ IT professionals help companies adopt and manage disruptive technologies to better compete

EXPERIENCED

11,800

Systems currently managed by Pythian

GLOBAL

400

Pythian experts in 35 countries

EXPERTS

2

Millennia of experience gathered and shared over 19 years

AGENDA

1. What to Expect
2. Motivation and Requirements
3. Ideas and The Solution
4. Oracle to Aurora (O2A)
5. Cutover
6. Issues we faced

What to Expect

WHAT TO EXPECT

Mind who’s in front of you

Oracle DB Expert *(wannabe)* - AWS Novice - MySQL Noob

WHAT TO EXPECT

My goal is not:

to teach you how to use AWS DMS

to provide the recipe of success with DMS

to dig deep in the technical details

My goal is:

to show what we did and how we did it

to have a fun discussion

to provide some hints on what to look out for

Motivation, Landscape and Requirements

The Motivation

DB Support contract expiring in 2018 Q3

Cost Optimization

Moving towards Cloud Native Future

Decision to move away from Oracle DBs

and use AWS Aurora MySQL for everything

I might be using terms «Aurora» and «MySQL» interchangeably

The Landscape

The project starts in Q1

6 months till the deadline

… and it is a hard deadline

THE LANDSCAPE

One RDS for Oracle instance (11.2)

Used by four applications

Can't move all apps at the same time:

Not enough resources

Too much risk

This presentation is only about Phase 1

Milestones:

Phase 1 – move one app and it's schema (400G)

Phase 2 – Move the 3 others later

The Landscape

Complications:

Cross schema FKs

Remaining apps still need access to data in Oracle for 2+ months

Good news:

Only read only access required

Not much code in the DB
few triggers / views and packages

The Requirements

Is that enough?

The First app to be moved in 4 months

1 practice switchover available – in QA

1 hour of downtime for the cutover

Not enough time for an offline move

Conclusions:

Replication or “Incremental refresh”

The Ideas and The Solution

THE IDEAS

Heterogenous Connectivity?

Connect Oracle to Aurora (like a DB link)

Refresh data on Aurora until the cutover

Use something similar to MVs

... or triggers to propagate DML

But…:

Oracle Database Gateway for MySQL not available on RDS

Sounds way too complicated anyway

The Ideas

Replication?

Replication is the way to limit the cutover time

Options:

Oracle Golden Gate was NOT considered

AWS Database Migration Service (DMS)

The Solution

AWS Database Migration Service was selected!

Advantages:

+ Built by AWS – well tested on RDS

+ One Stop Shop (AWS)

+ Supports CDC

Where have we heard this saying?

Doubts:

Previous experience from a year ago

when the tool was just recently announced was not great

The Solution

How DMS works:

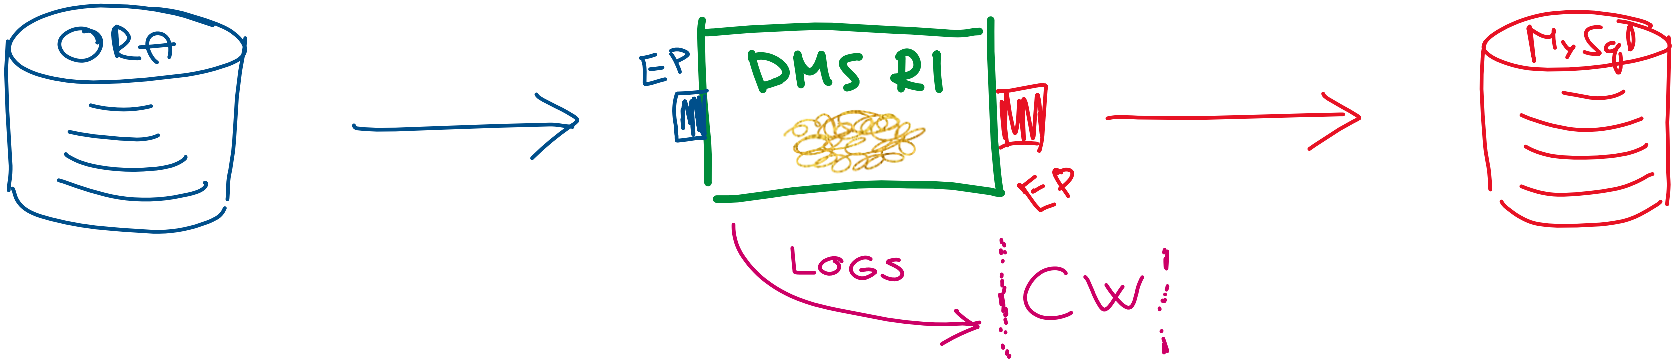

Full Load: Table by Table in parallel

Ongoing Replication: CDC, Logminer based

THE SOLUTION: REPLICATION APPROACH A

Before the Cutover of Application 1

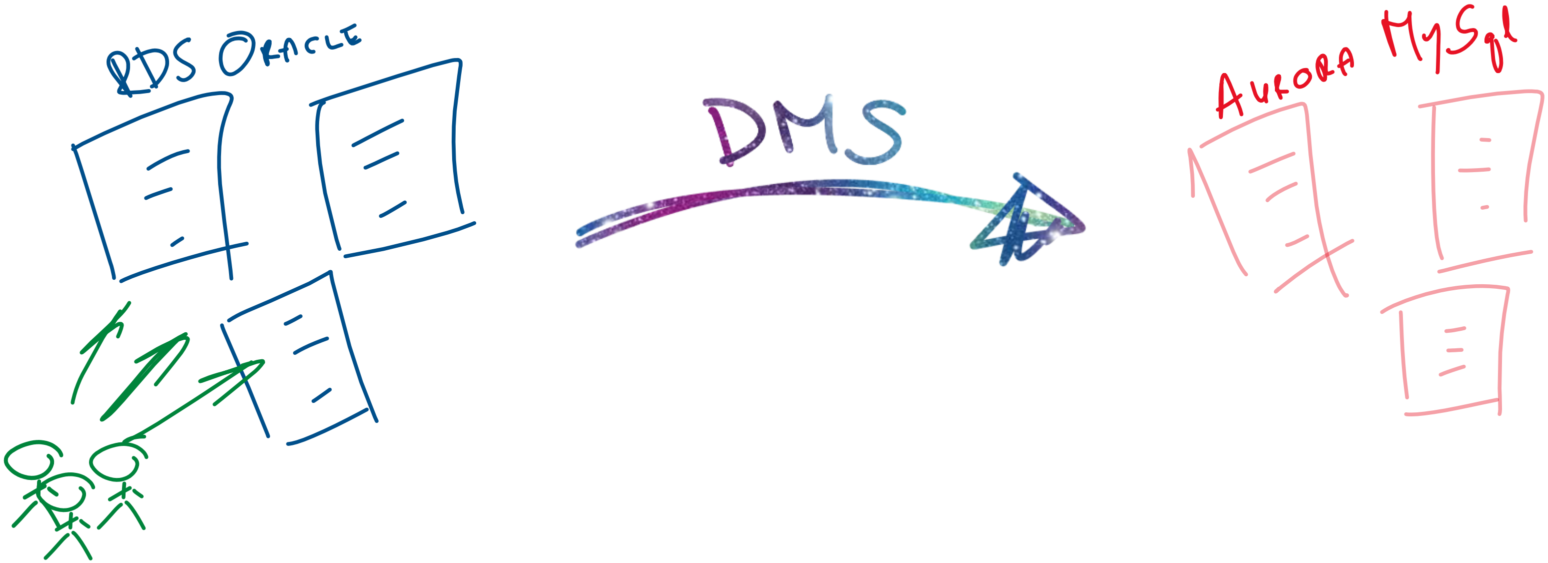

The Solution: Replication Approach A

After the Cutover of Application 1

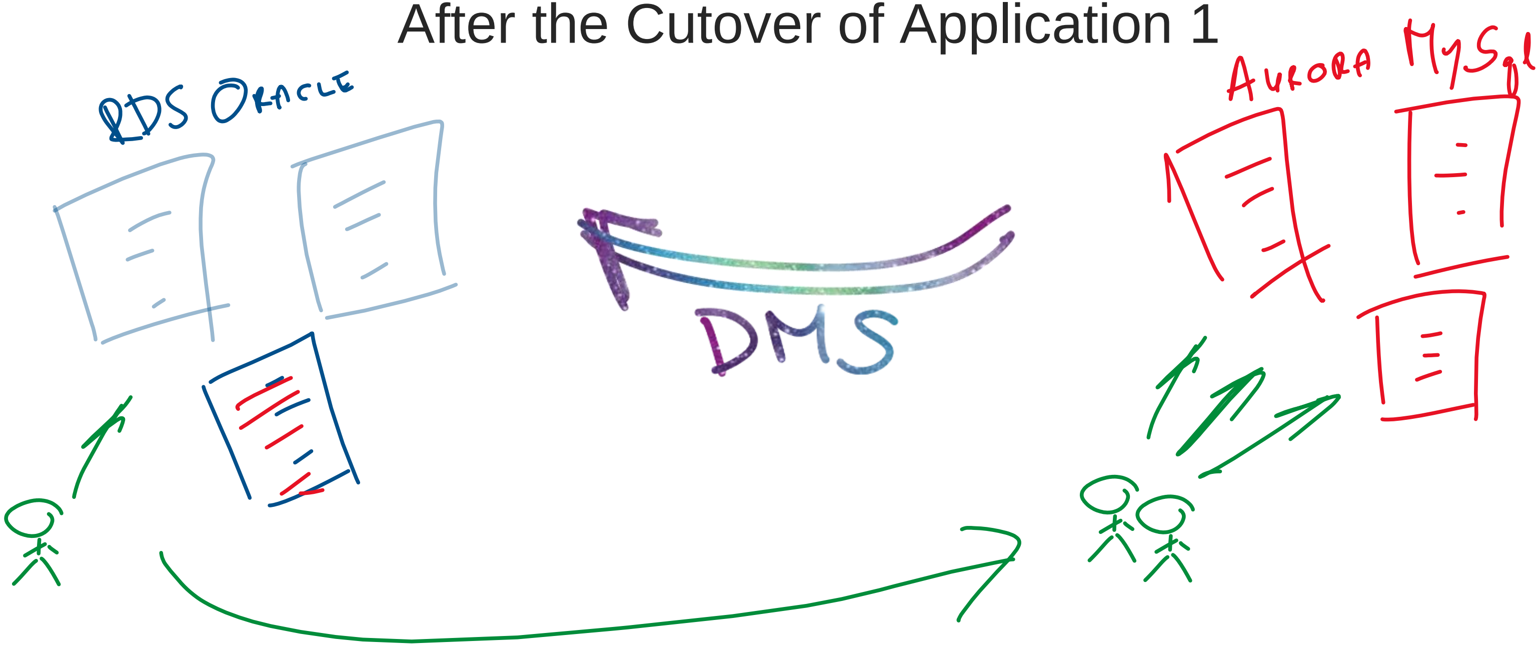

The Solution: Replication Approach A

What happens at the cutover:

1. Stop the app
2. Stop the DMS replication O2A
3. Reconnect the app to Aurora MySQL
4. Start replication from Aurora to Oracle (A2O)

... But there’s a problem ...

The Solution: Replication Approach A

What happens at the cutover:

1. Stop the app
2. Stop the DMS replication O2A
3. Reconnect the app to Aurora MySQL
4. Start replication from Aurora to Oracle (A2O)

The problem:

Trusting the A2O replication will work correctly

Can’t validate the data for A2O before it’s gone live

The Solution: Replication Approach B

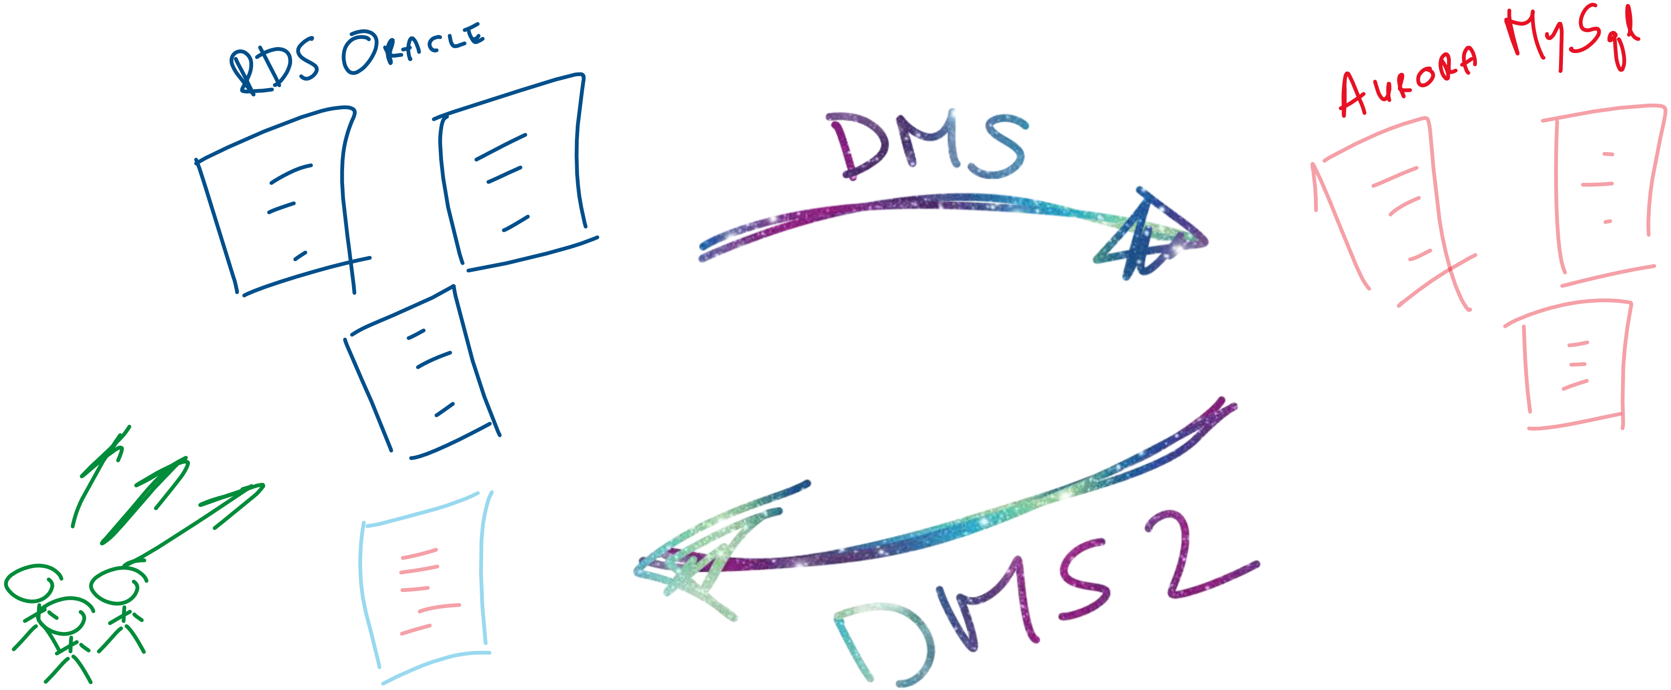

The Solution: Replication Approach B

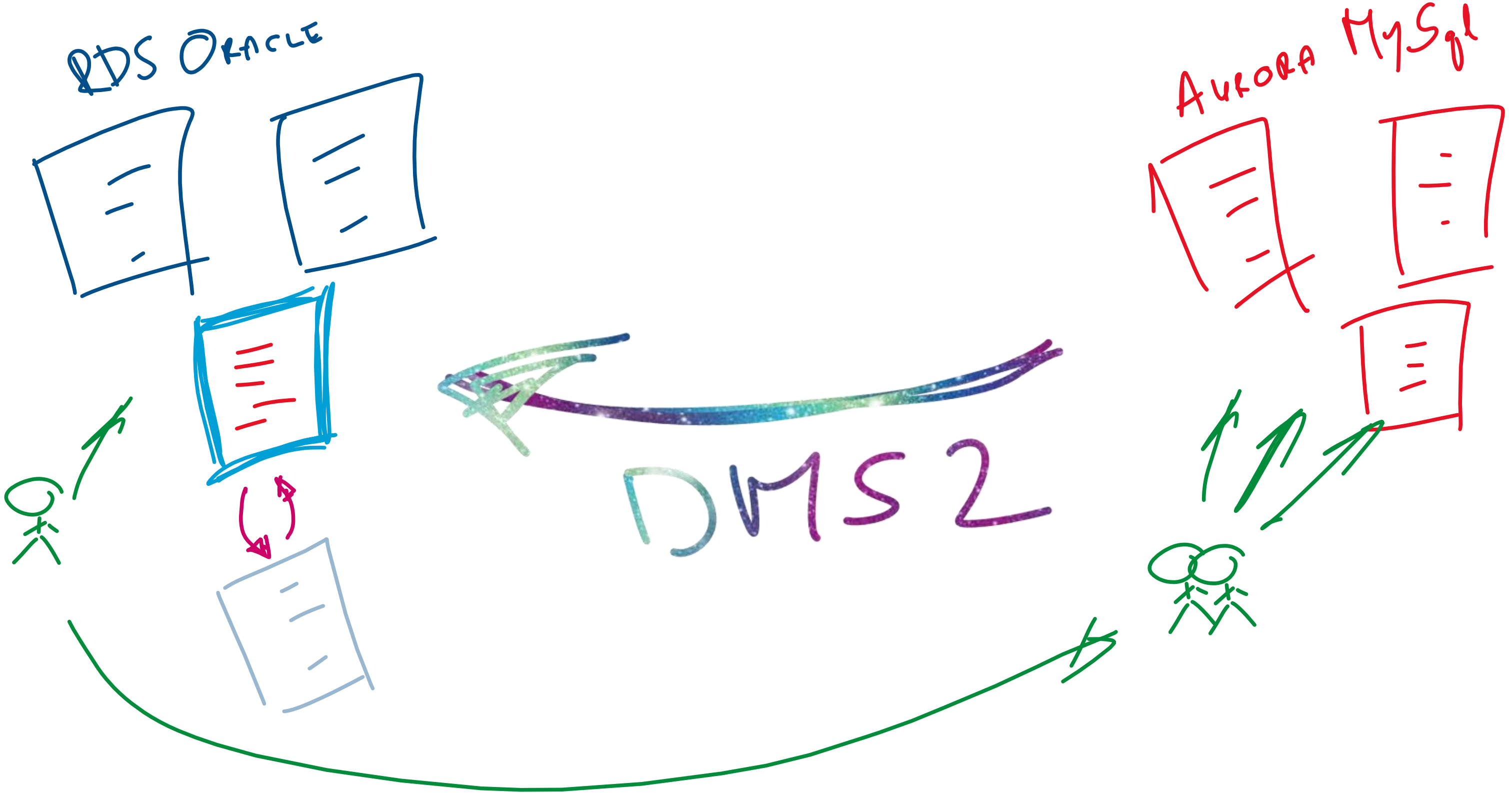

The Solution: Replication Approach B

Both Replication tasks running at the same time:

- Perform data validations
- Assess the stability of the replication task
- Observe logs for issues
- Prepare for the cutover

ORACLE TO AURORA (O2A)

O2A: Learning Curve

Lots of documentation available:

AWS Database Migration Service Documentation

Migrating an On-Premises Oracle Database to Amazon Aurora MySQL

Migrating an Amazon RDS Oracle Database to Amazon Aurora MySQL

Using an Oracle Database as a Source for AWS DMS

Using a MySQL-Compatible Database as a Target for AWS Database Migration Service

Pro Tip: Read the documentation before you start the work ;)

O2A: Overview of Implementation

1. Create target Aurora MySQL instances
2. Install AWS Schema Conversion Tool (SCT) on Your Computer
3. Test the connectivity to Oracle and Aurora MySQL
4. Use SCT to convert the schema from Oracle to MySQL
5. Validate the Schema - compare objects between DBs
6. Create a AWS DMS Replication Instance
7. Create AWS DMS Source and Target Endpoints
8. Create and run the AWS DMS Replication Task(s)
9. Verify the Data Migration is successful
10. Remove the Unnecessary Resources

Where is the cutover?

O2A: 1. Create target Aurora instances

It’s tempting to be agile, slow down!

Behavior differences between Oracle and Aurora MySQL

Think about the final settings you’ll need to have in the end

Is
Š == S == Ŝ
?

Character sets (`utf8` vs. `utf8mb4`)

🖕
Bytes: F0 9F
96 95

Collation (`utf8_bin` vs. `utf8_general_ci`)

Case insensitive table names (`lower_case_table_names`)

Slow query logging (`long_query_time` and `slow_query_log`)

«Performance
insights»
might help too

Defaults to 0
(NO)

O2A: 2. Install AWS Schema Conversion Tool (SCT) on Your Computer

Supported on Windows, Mac OS X, Ubuntu and Fedora

Install required Database Drivers:

- Oracle Database 12.1.0.1 JDBC Driver
- JDBC Driver for MySQL (Connector/J)

There's nothing really exciting about this part of work ☺

O2A: 3. Test the connectivity to Oracle and Aurora MySQL

Connectivty setup depends on your «landscape».

Controlled VDI in our case

... with only ssh access

... to an EC2 bastion host (a jumpbox)

... that had access to the DBs!

Few ssh-tunnels later
I was connected to both DBs from the SCT.

O2A: 4. USE SCT TO CONVERT THE SCHEMA FROM ORACLE TO MYSQL

Database Migration Assessment Report (PDF)

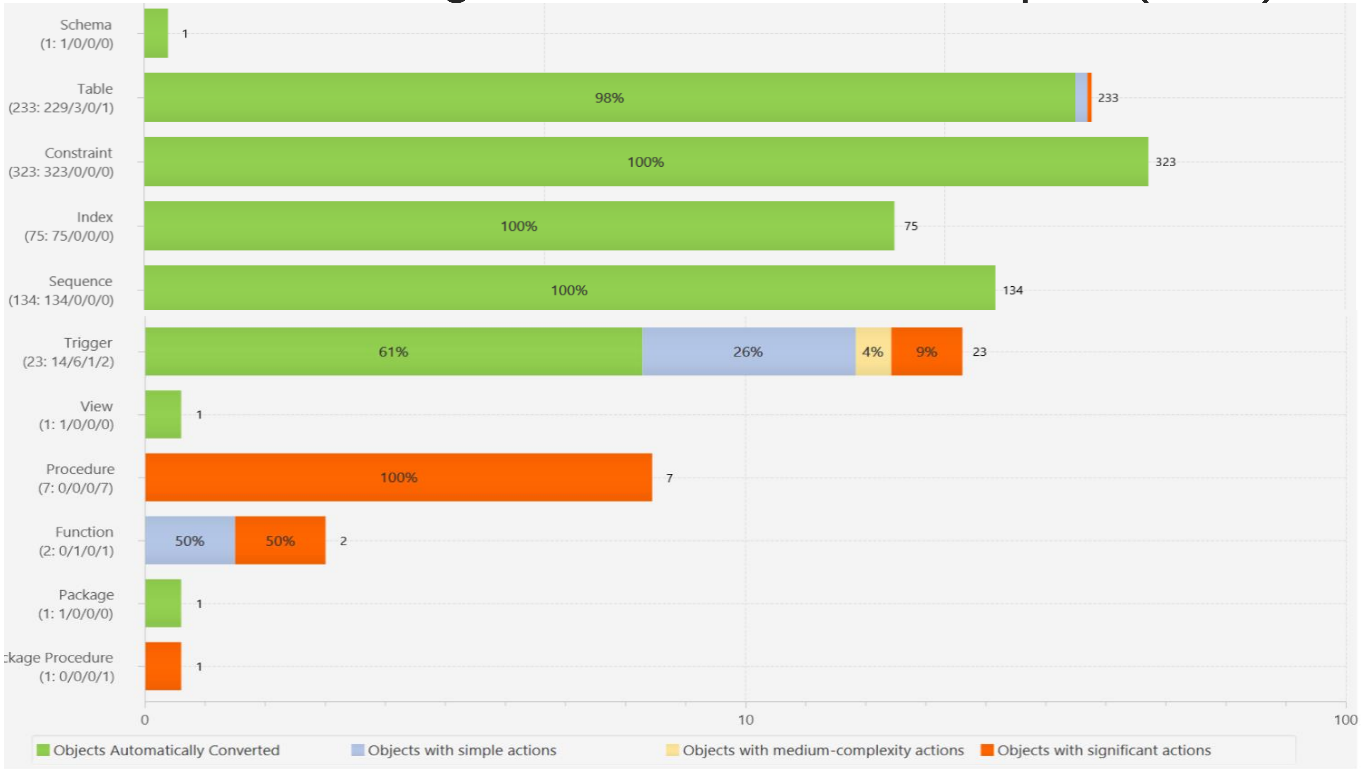

O2A: 4. Use SCT to convert the schema from Oracle to MySQL

- Issue 102: MySQL doesn't support the MERGE statement
- Issue 198: MySQL doesn't support GLOBAL TEMPORARY TABLE
- Issue 207: MySQL doesn't support function indexes
- Issue 326: MySQL doesn't support constraints with the status DISABLED
- Issue 332: MySQL doesn't support the procedure dbms_output.put_line
- Issue 340: Unable to convert functions
- Issue 345: The handler might cover not all cases
- Issue 350: The function cannot use statements that explicitly or implicitly begin or end a transaction, such as START TRANSACTION, COMMIT, or ROLLBACK
- Issue 524: MySQL doesn't support triggers for multiple events
- Issue 588: MySQL doesn't support more than one trigger per event. A few triggers were merged
- Issue 9996: Internal Converter error occurred

O2A: 4. Use SCT to convert the schema from Oracle to MySQL

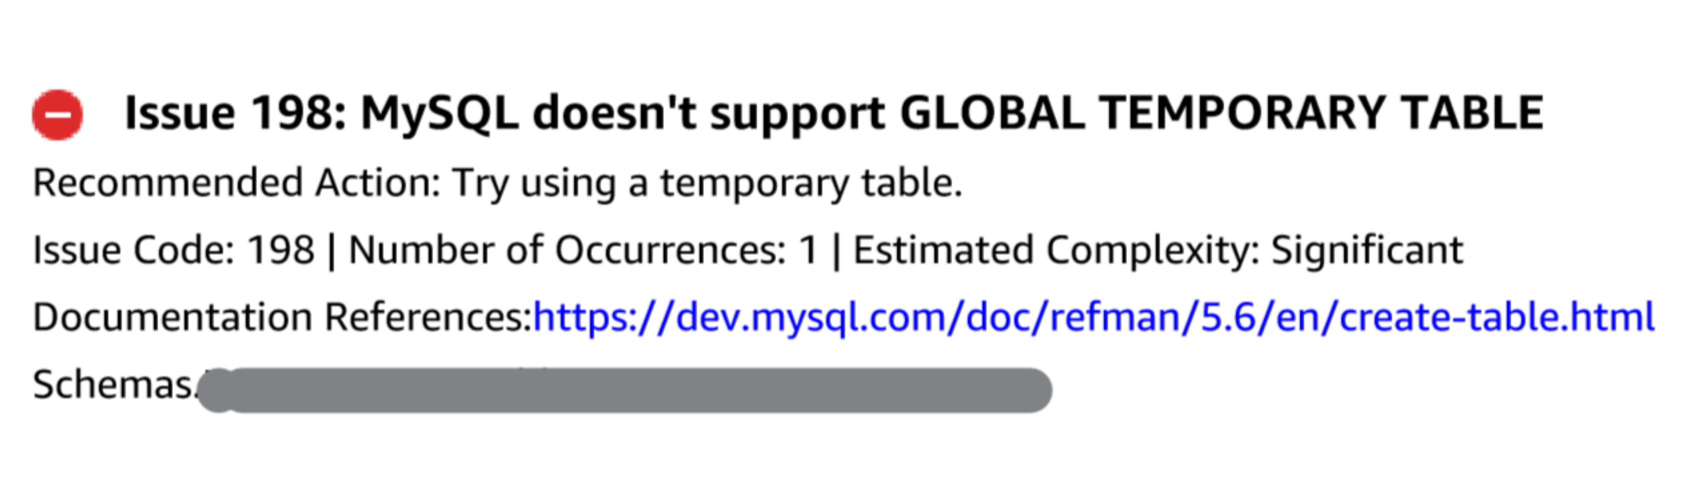

Issue 198: MySQL doesn't support GLOBAL TEMPORARY TABLE

Recommended Action: Try using a temporary table.

Issue Code: 198 | Number of Occurrences: 1 | Estimated Complexity: Significant

Documentation References:https://dev.mysql.com/doc/refman/5.6/en/create-table.html

Schemas:[redacted]

O2A: 4. Use SCT to convert the schema from Oracle to MySQL

Generate the Schema Creation Script

The SCT can:

– Create a «Project» file

– 1 to 1 source-target object mappings

– Record the automatically converted target object definitions

– Record the manually adjusted target object definitions

– Execute the create steps in the target database

Trust required
☺

Our way: Just create the initial script

– Take it outside SCT

– Managed it like any other source code

– wanted more control over datatypes

I.e. Standardise the *INT usage, dates vs timestamps

O2A: 4. Use SCT to convert the schema from Oracle to MySQL

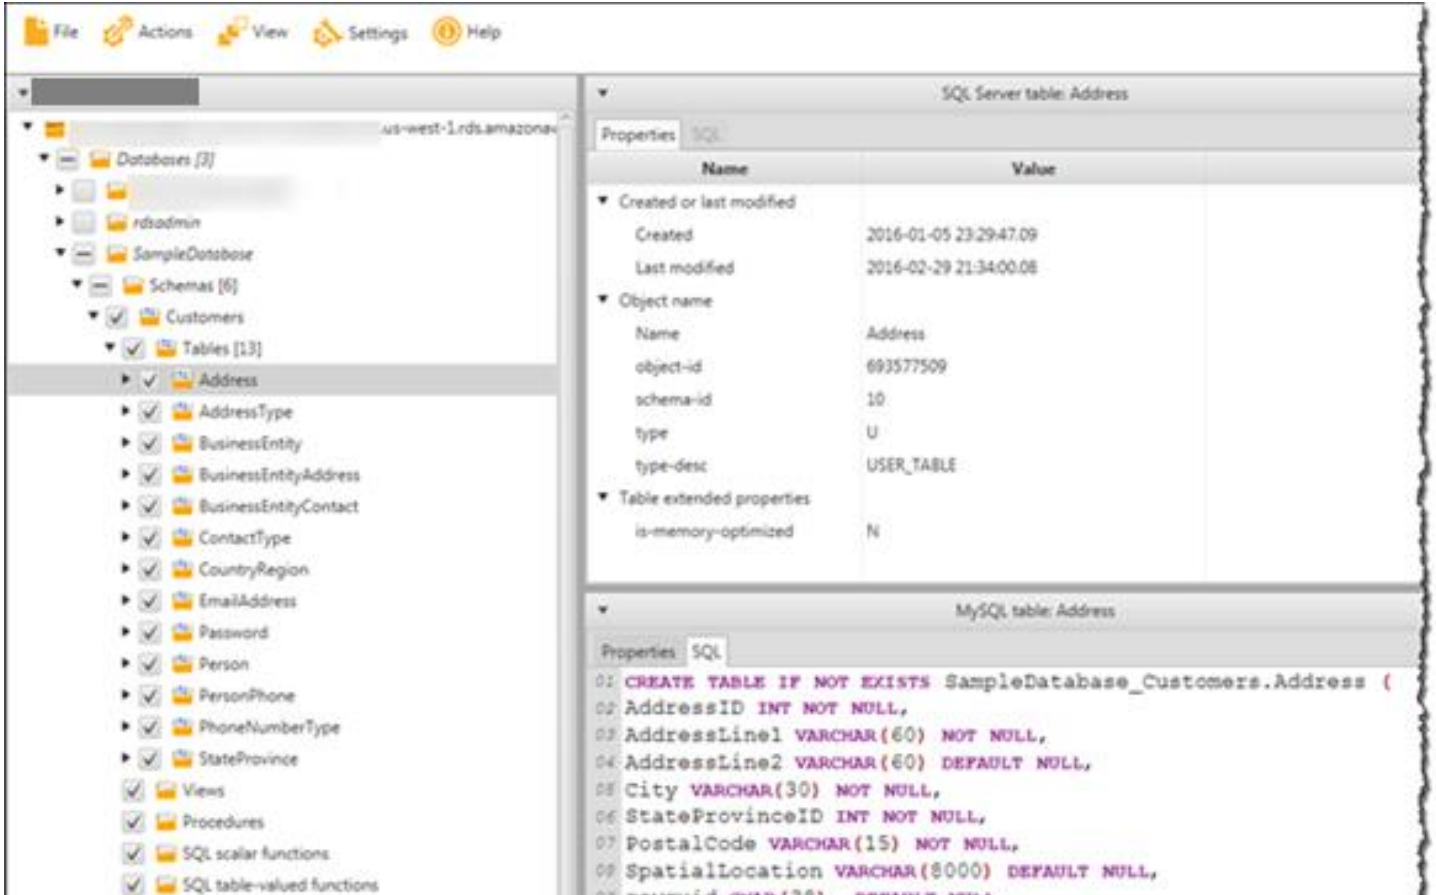

[redacted]

- [redacted] us-west-1.rds.amazona
  - Databases [3]
    - [redacted]
    - rdsadmin
    - SampleDatabase
      - Schemas [6]
        - ☑ Customers
          - ☑ Tables [13]
            - ☑ Address
            - ☑ AddressType
            - ☑ BusinessEntity
            - ☑ BusinessEntityAddress
            - ☑ BusinessEntityContact
            - ☑ ContactType
            - ☑ CountryRegion
            - ☑ EmailAddress
            - ☑ Password
            - ☑ Person
            - ☑ PersonPhone
            - ☑ PhoneNumberType
            - ☑ StateProvince
          - ☑ Views
          - ☑ Procedures
          - ☑ SQL scalar functions
          - ☑ SQL table-valued functions

SQL Server table: Address

Properties | SQL

Name	Value
Created or last modified	
Created	2016-01-05 23:29:47.09
Last modified	2016-02-29 21:34:00.08
Object name	
Name	Address
object-id	693577509
schema-id	10
type	U
type-desc	USER_TABLE
Table extended properties	
is-memory-optimized	N

MySQL table: Address

Properties | SQL

```
CREATE TABLE IF NOT EXISTS SampleDatabase_Customers.Address (
AddressID INT NOT NULL,
AddressLine1 VARCHAR(60) NOT NULL,
AddressLine2 VARCHAR(60) DEFAULT NULL,
City VARCHAR(30) NOT NULL,
StateProvinceID INT NOT NULL,
PostalCode VARCHAR(15) NOT NULL,
SpatialLocation VARCHAR(8000) DEFAULT NULL,
```

O2A: 4. Use SCT to convert the schema from Oracle to MySQL

SCT helped a lot!

Majority of objects were converted automatically without issues

Seeing a ready-made script was helpful

Remember,
I’m a
MySQL noob

O2A: 5. Validate the Schema – compare objects between DBs

Extract the object counts from both DBs and compare

The process we ran made sure it was OK without this comparison

Several iterations of:

– Drop the target DB

– Run the schema creation script

– Check for errors

May be more thorough validation is required

when everything's done with SCT?

O2A: 6. Create an AWS DMS Replication Instance

It's basically just a pre-configured EC2 created by a Wizard by providing parameters:

- Name
- Instance type (vCPU / RAM)
- Storage size
- Subnet / VPC Security Groups
- KMS Master Key
- ...
- Extra Connetion Attributes

O2A: 7. Create AWS DMS Source and Target Endpoints

Endpoints:

Source / Target

connection description + special attributes

Replication instance → DB Server

Connection test

cNAME issue with writer and the read replicas

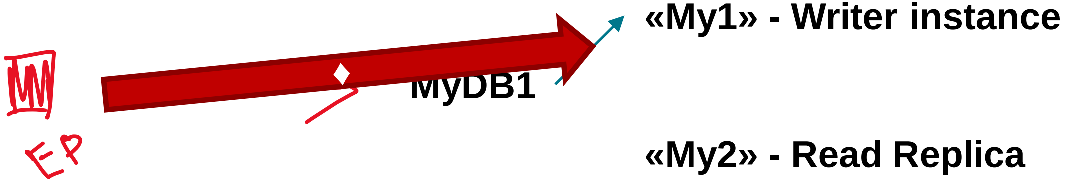

«My1» - Writer instance

«My2» - Read Replica

O2A: 7. Source Endpoint – special attributes

```
aws dms describe-endpoints
..
{
"Username": "pythian",
"Status": "active",
"EndpointArn": "xxx",
"ServerName": "yyy.vvv.us-east-1.rds.amazonaws.com",
"EndpointType": "SOURCE",
"SslMode": "none",
"KmsKeyId": "***",
"ExtraConnectionAttributes": "addSupplementalLogging=Y",
"DatabaseName": "PROD",
"EngineDisplayName": "Oracle",
"EngineName": "oracle",
"EndpointIdentifier": «source_db01",
"Port": 1521
},
```

O2A: 7. TARGET ENDPOINT - SPECIAL ATTRIBUTES

```
aws dms describe-endpoints
..
{
"Username": "pythian",
"Status": "active",
"EndpointArn": "xxx",
"ServerName": "xxxxxxxxxxxxxxx.yyyyyyyyyyy.us-east-1.rds.amazonaws.com",
"EndpointType": "TARGET",
"SslMode": "none",
"KmsKeyId": "***",
"ExtraConnectionAttributes": "initstmt=SET FOREIGN_KEY_CHECKS=0;parallelLoadThreads=1;afterConnectScript=set
session sql_mode =
'STRICT_TRANS_TABLES,NO_ZERO_IN_DATE,NO_ZERO_DATE,ERROR_FOR_DIVISION_BY_ZERO,NO_AUTO_CREATE_USER,NO_ENGINE_SUBST
ITUTION,NO_AUTO_VALUE_ON_ZERO,STRICT_ALL_TABLES'",
"EngineDisplayName": "Amazon Aurora MySQL",
"EngineName": "aurora",
"EndpointIdentifier": "target-db01",
"Port": 3306
},
..
```

O2A: 8. CREATE AND RUN THE AWS DMS REPLICATION TASK(S)

GUI-way ... and it’s troubles

Simple, unless you need a bit of extra

Lots of tables / rules

Nondefault Logging

Specific settings

Stop → Start/Restart

O2A: 8. DMS TAK GUI

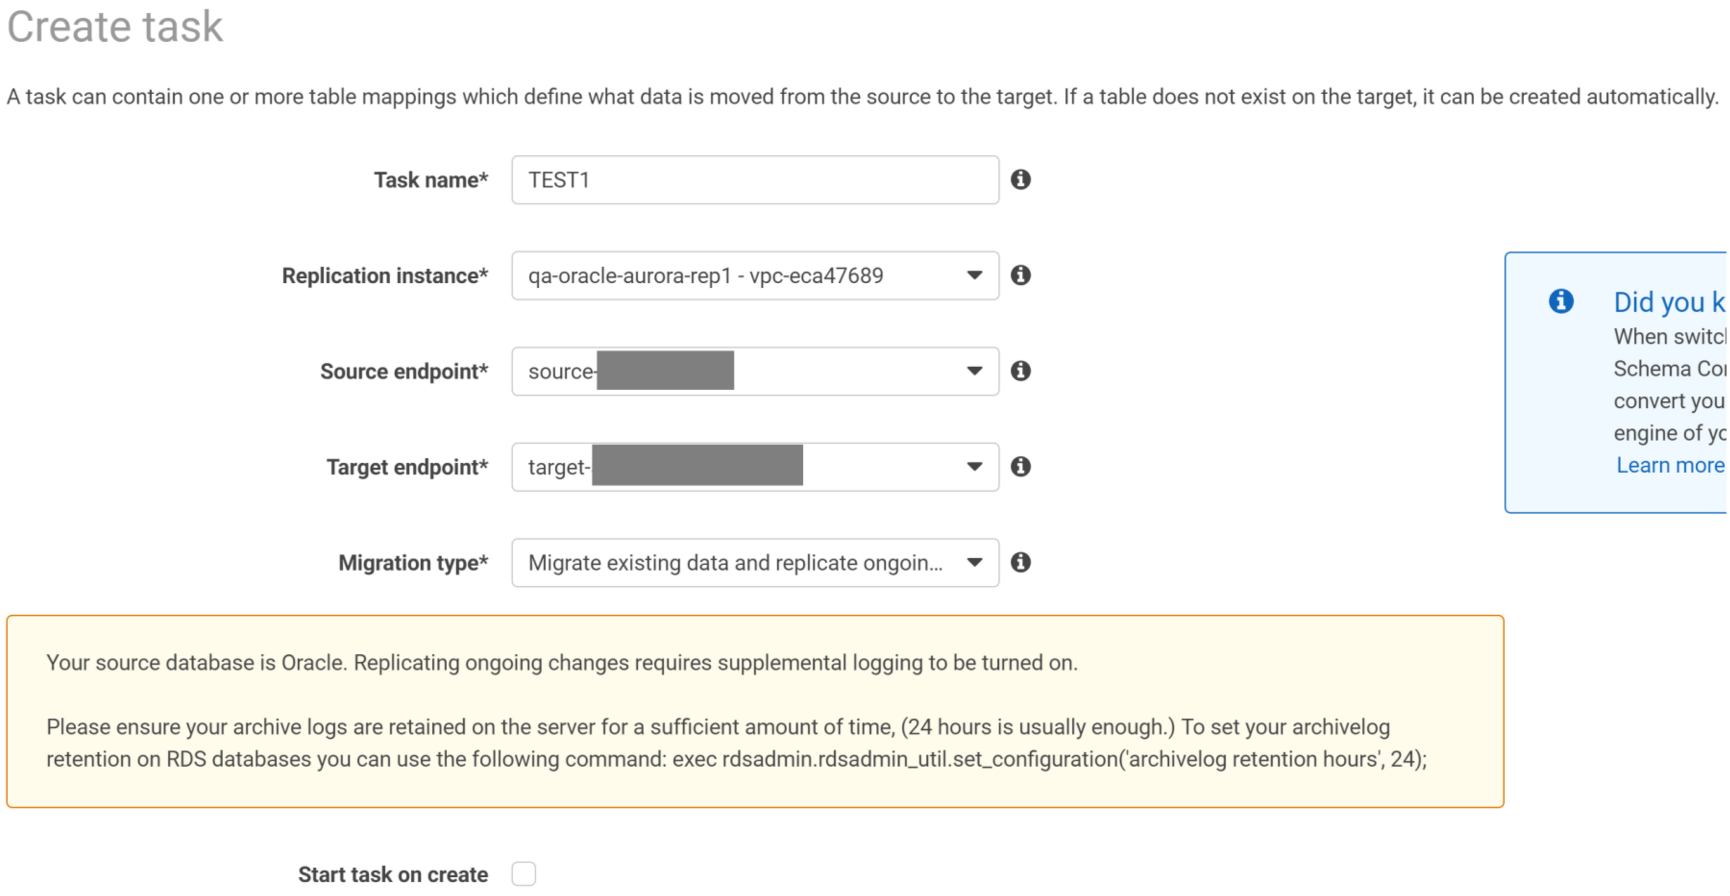

Create task

A task can contain one or more table mappings which define what data is moved from the source to the target. If a table does not exist on the target, it can be created automatically.

Task name*	TEST1
Replication instance*	qa-oracle-aurora-rep1 - vpc-eca47689
Source endpoint*	source-[redacted]
Target endpoint*	target-[redacted]
Migration type*	Migrate existing data and replicate ongoin...

Did you k
When switc
Schema Co
convert you
engine of yo
Learn more

Your source database is Oracle. Replicating ongoing changes requires supplemental logging to be turned on.

Please ensure your archive logs are retained on the server for a sufficient amount of time, (24 hours is usually enough.) To set your archivelog retention on RDS databases you can use the following command: exec rdsadmin.rdsadmin_util.set_configuration('archivelog retention hours', 24);

Start task on create □

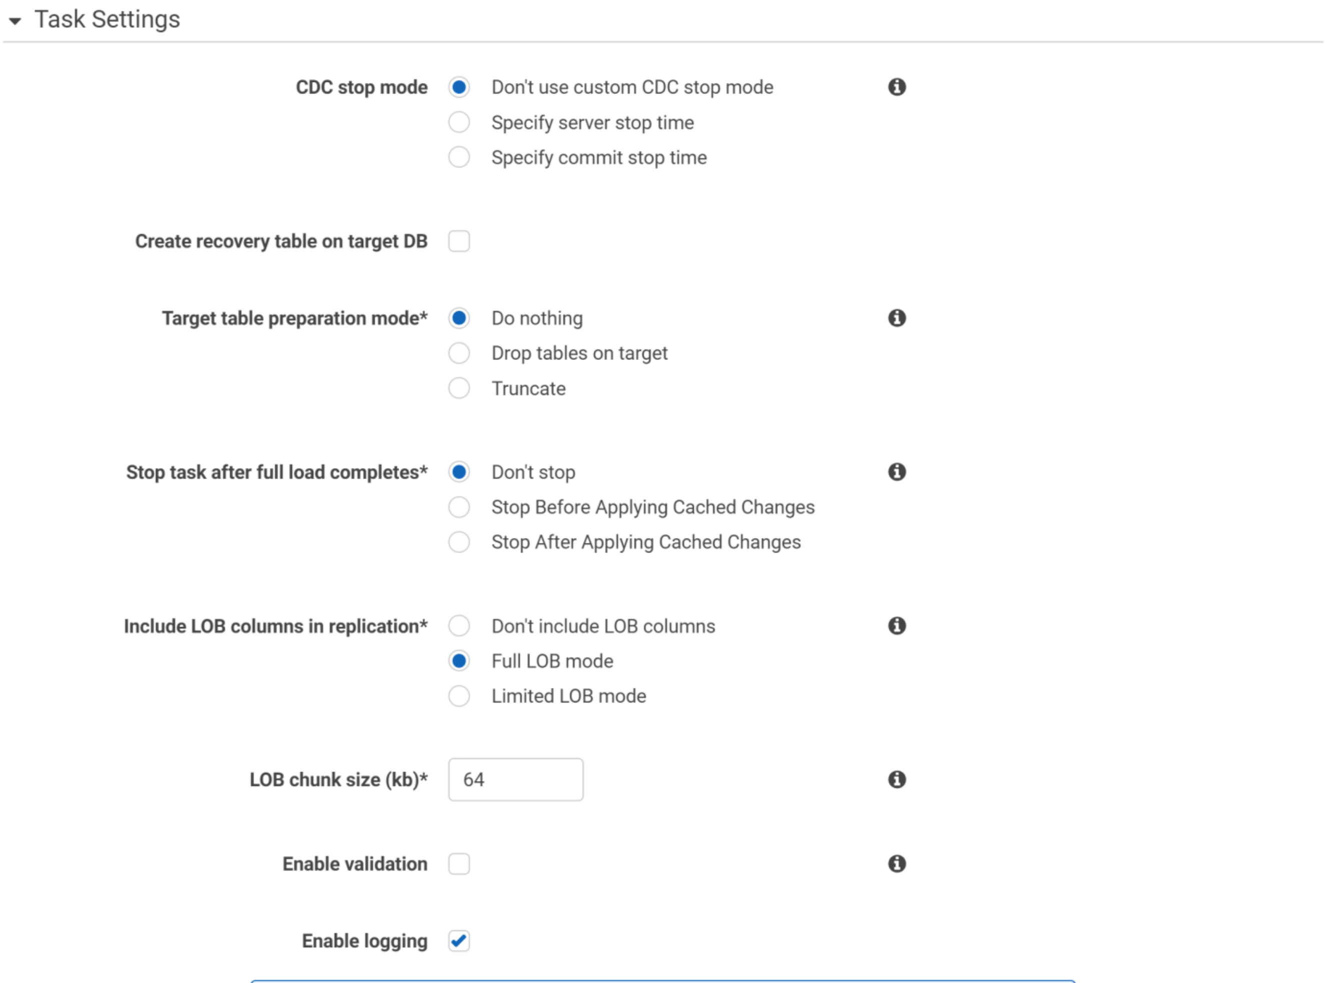

▾ Task Settings

CDC stop mode
- ● Don't use custom CDC stop mode
- ○ Specify server stop time
- ○ Specify commit stop time

Create recovery table on target DB ☐

Target table preparation mode*
- ● Do nothing
- ○ Drop tables on target
- ○ Truncate

Stop task after full load completes*
- ● Don't stop
- ○ Stop Before Applying Cached Changes
- ○ Stop After Applying Cached Changes

Include LOB columns in replication*
- ○ Don't include LOB columns
- ● Full LOB mode
- ○ Limited LOB mode

LOB chunk size (kb)* 64

Enable validation ☐

Enable logging ☑

O2A: 8. DMS Tak GUI

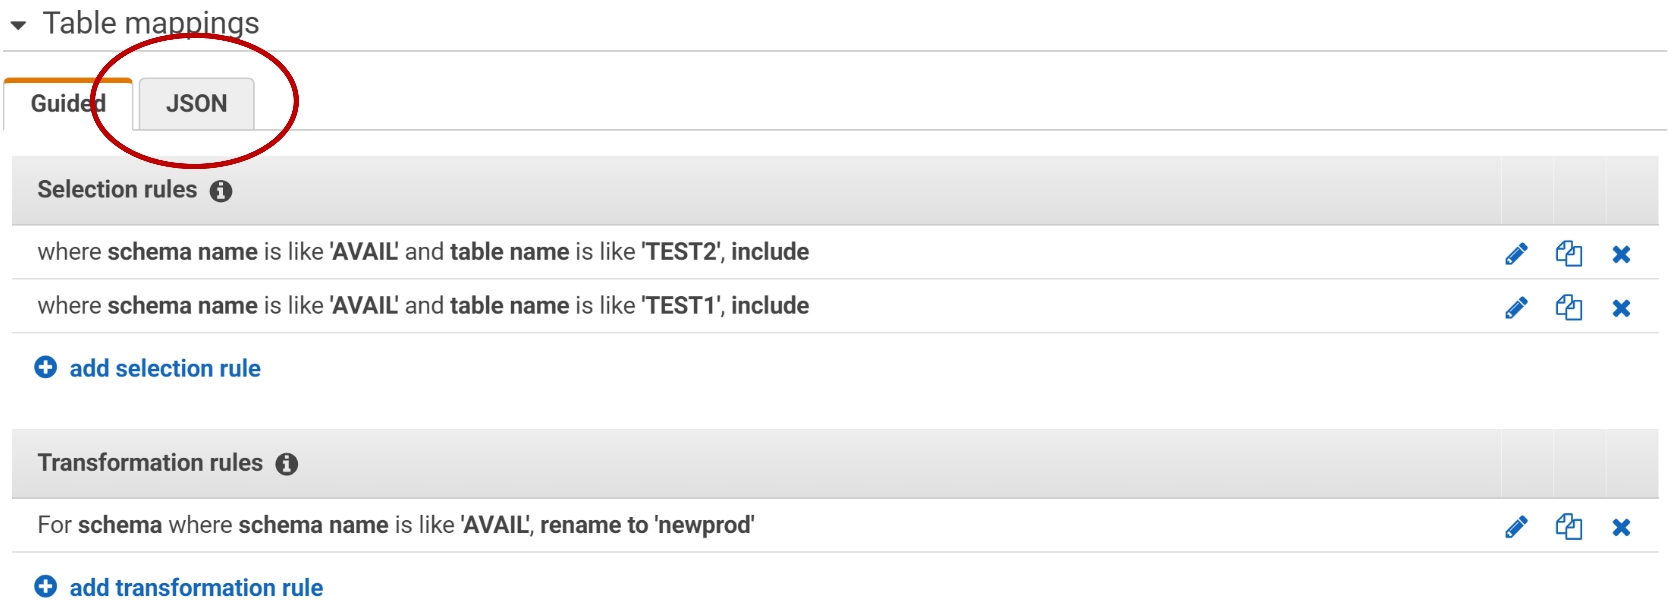

Table mappings

Selection rules

where **schema name** is like **'AVAIL'** and **table name** is like **'TEST2'**, **include**

where **schema name** is like **'AVAIL'** and **table name** is like **'TEST1'**, **include**

Transformation rules

For **schema** where **schema name** is like **'AVAIL'**, **rename to 'newprod'**

O2A: 8. Create and run the AWS DMS Replication Task(s)

```
{
  "rules": [
    {
      "rule-type": "selection",
      "rule-id": "1",
      "rule-name": "1",
      "object-locator": {
        "schema-name": "AVAIL",
        "table-name": "TEST2"
      },
      "rule-action": "include"
    },
...
```

```
    {
      "rule-type": "transformation",
      "rule-id": «75",
      "rule-name": «75",
      "rule-target": "schema",
      "object-locator": {
        "schema-name": "OLD_PROD"
      },
      "rule-action": "rename",
      "value": "newprod"
    }
  ]
}
```

O2A: 8. Create and run the AWS DMS Replication

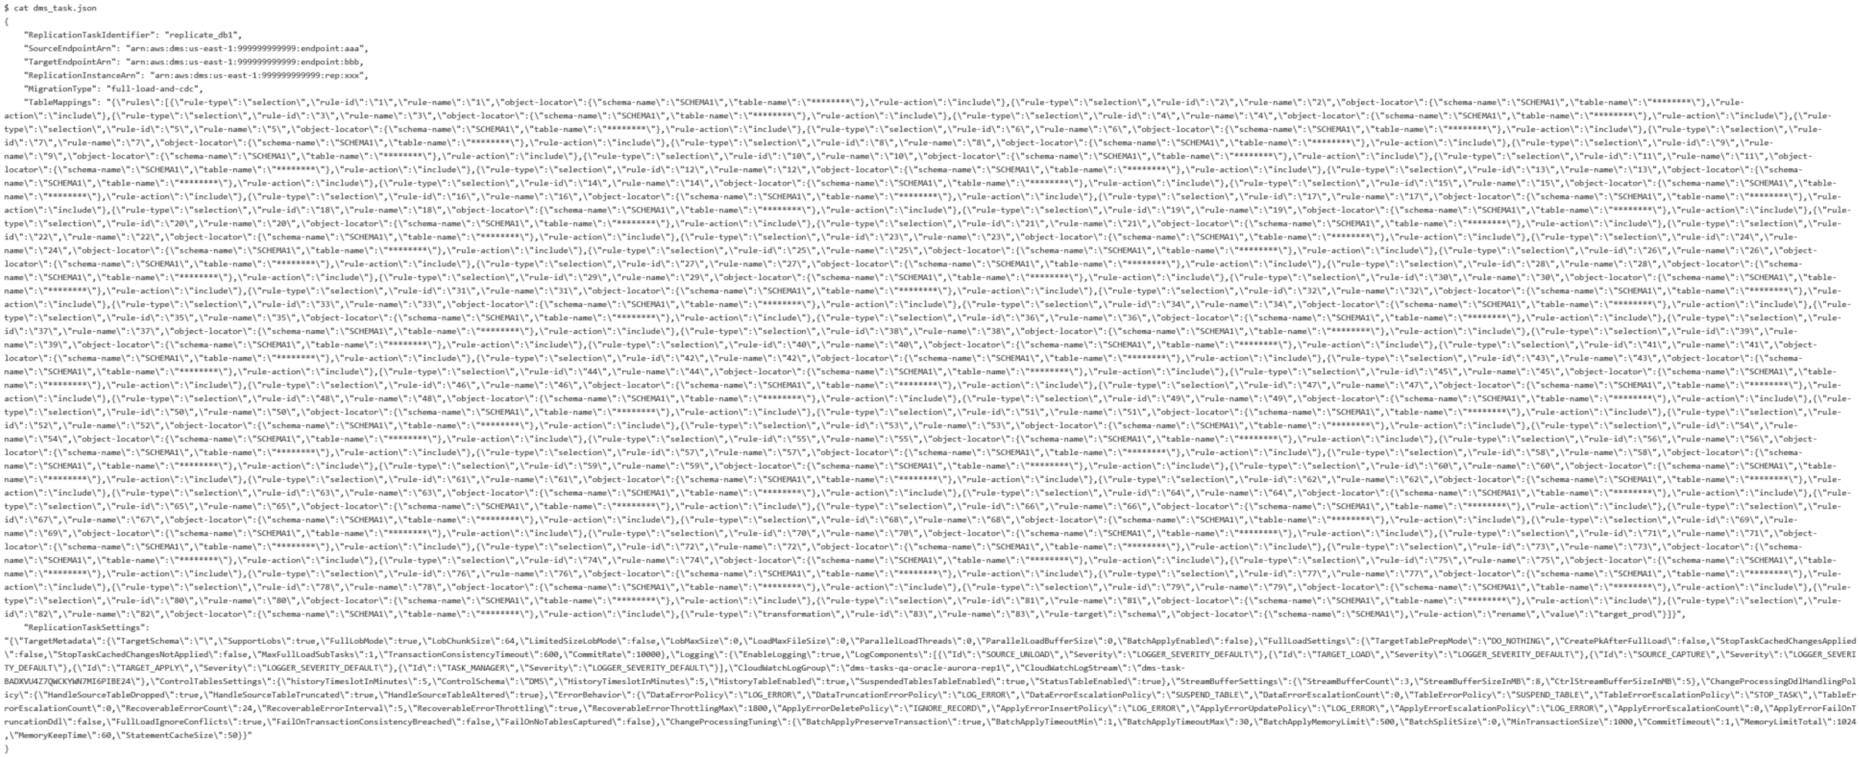

```
$ cat dms_task.json
{
    "ReplicationTaskIdentifier": "replicate_db1",
    "SourceEndpointArn": "arn:aws:dms:us-east-1:999999999999:endpoint:aaa",
    "TargetEndpointArn": "arn:aws:dms:us-east-1:999999999999:endpoint:bbb,
    "ReplicationInstanceArn": "arn:aws:dms:us-east-1:999999999999:rep:xxx",
    "MigrationType": "full-load-and-cdc",
    "TableMappings": "{\"rules\":[{\"rule-type\":\"selection\",\"rule-id\":\"1\",\"rule-name\":\"1\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"2\",\"rule-name\":\"2\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-
action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"3\",\"rule-name\":\"3\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"4\",\"rule-name\":\"4\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-
type\":\"selection\",\"rule-id\":\"5\",\"rule-name\":\"5\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"6\",\"rule-name\":\"6\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-
id\":\"7\",\"rule-name\":\"7\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"8\",\"rule-name\":\"8\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"9\",\"rule-
name\":\"9\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"10\",\"rule-name\":\"10\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"11\",\"rule-name\":\"11\",\"object-
locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"12\",\"rule-name\":\"12\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"13\",\"rule-name\":\"13\",\"object-locator\":{\"schema-
name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"14\",\"rule-name\":\"14\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"15\",\"rule-name\":\"15\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-
name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"16\",\"rule-name\":\"16\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"17\",\"rule-name\":\"17\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-
action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"18\",\"rule-name\":\"18\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"19\",\"rule-name\":\"19\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-
type\":\"selection\",\"rule-id\":\"20\",\"rule-name\":\"20\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"21\",\"rule-name\":\"21\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-
id\":\"22\",\"rule-name\":\"22\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"23\",\"rule-name\":\"23\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"24\",\"rule-
name\":\"24\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"25\",\"rule-name\":\"25\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"26\",\"rule-name\":\"26\",\"object-
locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"27\",\"rule-name\":\"27\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"28\",\"rule-name\":\"28\",\"object-locator\":{\"schema-
name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"29\",\"rule-name\":\"29\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"30\",\"rule-name\":\"30\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-
name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"31\",\"rule-name\":\"31\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"32\",\"rule-name\":\"32\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-
action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"33\",\"rule-name\":\"33\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"34\",\"rule-name\":\"34\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-
type\":\"selection\",\"rule-id\":\"35\",\"rule-name\":\"35\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"36\",\"rule-name\":\"36\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-
id\":\"37\",\"rule-name\":\"37\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"38\",\"rule-name\":\"38\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"39\",\"rule-
name\":\"39\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"40\",\"rule-name\":\"40\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"41\",\"rule-name\":\"41\",\"object-
locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"42\",\"rule-name\":\"42\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"43\",\"rule-name\":\"43\",\"object-locator\":{\"schema-
name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"44\",\"rule-name\":\"44\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"45\",\"rule-name\":\"45\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-
name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"46\",\"rule-name\":\"46\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"47\",\"rule-name\":\"47\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-
action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"48\",\"rule-name\":\"48\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"49\",\"rule-name\":\"49\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-
type\":\"selection\",\"rule-id\":\"50\",\"rule-name\":\"50\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"51\",\"rule-name\":\"51\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-
id\":\"52\",\"rule-name\":\"52\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"53\",\"rule-name\":\"53\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"54\",\"rule-
name\":\"54\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"55\",\"rule-name\":\"55\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"56\",\"rule-name\":\"56\",\"object-
locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"57\",\"rule-name\":\"57\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"58\",\"rule-name\":\"58\",\"object-locator\":{\"schema-
name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"59\",\"rule-name\":\"59\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"60\",\"rule-name\":\"60\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-
name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"61\",\"rule-name\":\"61\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"62\",\"rule-name\":\"62\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-
action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"63\",\"rule-name\":\"63\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"64\",\"rule-name\":\"64\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-
type\":\"selection\",\"rule-id\":\"65\",\"rule-name\":\"65\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"66\",\"rule-name\":\"66\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-
id\":\"67\",\"rule-name\":\"67\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"68\",\"rule-name\":\"68\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"69\",\"rule-
name\":\"69\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"70\",\"rule-name\":\"70\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"71\",\"rule-name\":\"71\",\"object-
locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"72\",\"rule-name\":\"72\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"73\",\"rule-name\":\"73\",\"object-locator\":{\"schema-
name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"74\",\"rule-name\":\"74\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"75\",\"rule-name\":\"75\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-
name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"76\",\"rule-name\":\"76\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"77\",\"rule-name\":\"77\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-
action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"78\",\"rule-name\":\"78\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"79\",\"rule-name\":\"79\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-
type\":\"selection\",\"rule-id\":\"80\",\"rule-name\":\"80\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-id\":\"81\",\"rule-name\":\"81\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"selection\",\"rule-
id\":\"82\",\"rule-name\":\"82\",\"object-locator\":{\"schema-name\":\"SCHEMA1\",\"table-name\":\"********\"},\"rule-action\":\"include\"},{\"rule-type\":\"transformation\",\"rule-id\":\"83\",\"rule-name\":\"83\",\"rule-target\":\"schema\",\"object-locator\":{\"schema-name\":\"SCHEMA1\"},\"rule-action\":\"rename\",\"value\":\"target_prod\"}]}",
    "ReplicationTaskSettings":
"{\"TargetMetadata\":{\"TargetSchema\":\"\",\"SupportLobs\":true,\"FullLobMode\":true,\"LobChunkSize\":64,\"LimitedSizeLobMode\":false,\"LobMaxSize\":0,\"LoadMaxFileSize\":0,\"ParallelLoadThreads\":0,\"ParallelLoadBufferSize\":0,\"BatchApplyEnabled\":false},\"FullLoadSettings\":{\"TargetTablePrepMode\":\"DO_NOTHING\",\"CreatePkAfterFullLoad\":false,\"StopTaskCachedChangesApplied
\":false,\"StopTaskCachedChangesNotApplied\":false,\"MaxFullLoadSubTasks\":1,\"TransactionConsistencyTimeout\":600,\"CommitRate\":10000},\"Logging\":{\"EnableLogging\":true,\"LogComponents\":[{\"Id\":\"SOURCE_UNLOAD\",\"Severity\":\"LOGGER_SEVERITY_DEFAULT\"},{\"Id\":\"TARGET_LOAD\",\"Severity\":\"LOGGER_SEVERITY_DEFAULT\"},{\"Id\":\"SOURCE_CAPTURE\",\"Severity\":\"LOGGER_SEVERI
TY_DEFAULT\"},{\"Id\":\"TARGET_APPLY\",\"Severity\":\"LOGGER_SEVERITY_DEFAULT\"},{\"Id\":\"TASK_MANAGER\",\"Severity\":\"LOGGER_SEVERITY_DEFAULT\"}],\"CloudWatchLogGroup\":\"dms-tasks-qa-oracle-aurora-rep1\",\"CloudWatchLogStream\":\"dms-task-
BADXVU4Z7QWCKYWN7MI6PIBE24\"},\"ControlTablesSettings\":{\"historyTimeslotInMinutes\":5,\"ControlSchema\":\"DMS\",\"HistoryTimeslotInMinutes\":5,\"HistoryTableEnabled\":true,\"SuspendedTablesTableEnabled\":true,\"StatusTableEnabled\":true},\"StreamBufferSettings\":{\"StreamBufferCount\":3,\"StreamBufferSizeInMB\":8,\"CtrlStreamBufferSizeInMB\":5},\"ChangeProcessingDdlHandlingPol
icy\":{\"HandleSourceTableDropped\":true,\"HandleSourceTableTruncated\":true,\"HandleSourceTableAltered\":true},\"ErrorBehavior\":{\"DataErrorPolicy\":\"LOG_ERROR\",\"DataTruncationErrorPolicy\":\"LOG_ERROR\",\"DataErrorEscalationPolicy\":\"SUSPEND_TABLE\",\"DataErrorEscalationCount\":0,\"TableErrorPolicy\":\"SUSPEND_TABLE\",\"TableErrorEscalationPolicy\":\"STOP_TASK\",\"TableEr
rorEscalationCount\":0,\"RecoverableErrorCount\":24,\"RecoverableErrorInterval\":5,\"RecoverableErrorThrottling\":true,\"RecoverableErrorThrottlingMax\":1800,\"ApplyErrorDeletePolicy\":\"IGNORE_RECORD\",\"ApplyErrorInsertPolicy\":\"LOG_ERROR\",\"ApplyErrorUpdatePolicy\":\"LOG_ERROR\",\"ApplyErrorEscalationPolicy\":\"LOG_ERROR\",\"ApplyErrorEscalationCount\":0,\"ApplyErrorFailOnT
runcationDdl\":false,\"FullLoadIgnoreConflicts\":true,\"FailOnTransactionConsistencyBreached\":false,\"FailOnNoTablesCaptured\":false},\"ChangeProcessingTuning\":{\"BatchApplyPreserveTransaction\":true,\"BatchApplyTimeoutMin\":1,\"BatchApplyTimeoutMax\":30,\"BatchApplyMemoryLimit\":500,\"BatchSplitSize\":0,\"MinTransactionSize\":1000,\"CommitTimeout\":1,\"MemoryLimitTotal\":1024
,\"MemoryKeepTime\":60,\"StatementCacheSize\":50}}"
}
```

O2A: 8. Create and run the AWS DMS Replication Task(s)

Some advanced options can't be changed from GUI

... Use AWS CLI and json definition of the task.

What we want to change?

"RecoverableErrorCount" = 24 (from -1)

LOGGER_SEVERITY_DETAILED_DEBUG

```
aws dms create-replication-task
--cli-input-json file://prod-ora2aurora.json
```

O2A: 9. Verify the Data Migration is successful

Time to stop the replication!

Is all data processed?

Is data migrated correctly?

Make sure the last redo sequence is processed!

DMS provides data validation!

Unfortunately:

Good for small, simple tables

Unreliable / Unpredictable for others

O2A: 9. Verify the Data Migration is successful

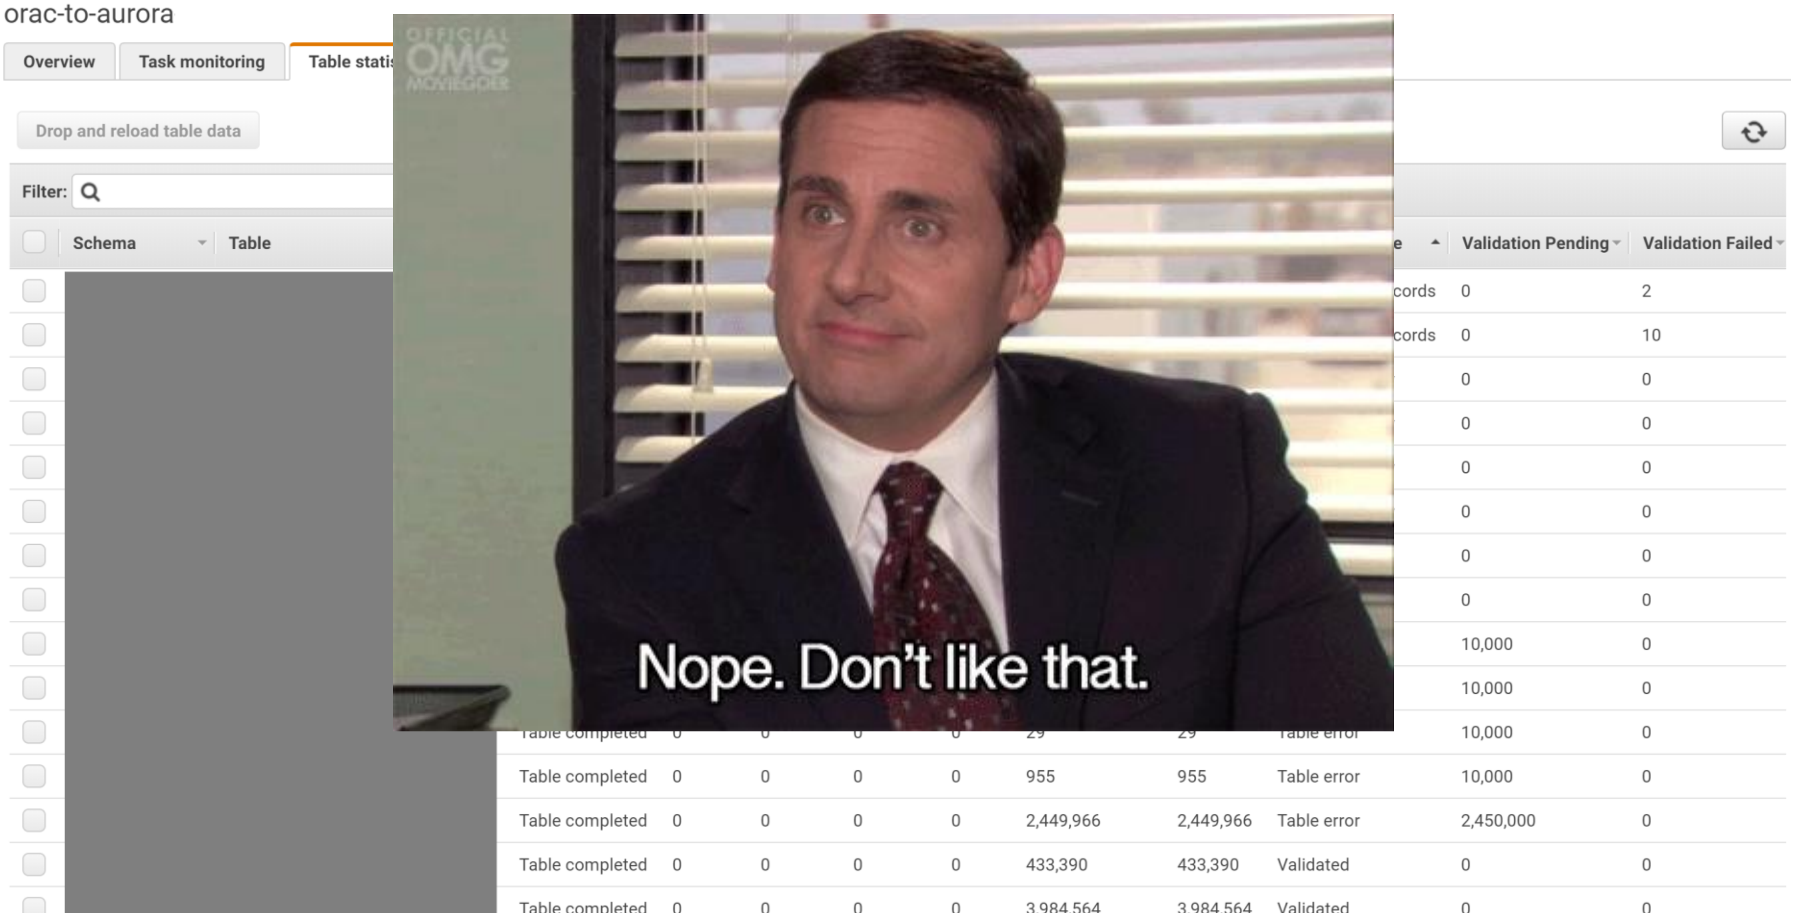

orac-to-aurora

□	Schema	Table								e ▲	Validation Pending	Validation Failed
□	[redacted]	[redacted]								cords	0	2
□	[redacted]	[redacted]								cords	0	10
□	[redacted]	[redacted]									0	0
□	[redacted]	[redacted]									0	0
□	[redacted]	[redacted]									0	0
□	[redacted]	[redacted]									0	0
□	[redacted]	[redacted]									0	0
□	[redacted]	[redacted]									0	0
□	[redacted]	[redacted]									10,000	0
□	[redacted]	[redacted]									10,000	0
□	[redacted]	[redacted]	Table completed	0	0	0	0	29	29	Table error	10,000	0
□	[redacted]	[redacted]	Table completed	0	0	0	0	955	955	Table error	10,000	0
□	[redacted]	[redacted]	Table completed	0	0	0	0	2,449,966	2,449,966	Table error	2,450,000	0
□	[redacted]	[redacted]	Table completed	0	0	0	0	433,390	433,390	Validated	0	0
□	[redacted]	[redacted]	Table completed	0	0	0	0	3,984,564	3,984,564	Validated	0	0

O2A: 9. DATA VALIDATION

Created a custom data extraction tool

- Scripts that generate data extraction queries
- Run queries against both databases
- Data is written to csv files
- «Difficult» columns are MD5-hashed
  – ones except numbers, dates and timestamps
- Compare the outputs with «diff»

O2A: 9. DATA VALIDATION

Data extraction in MySql

All numbers: DOUBLE, DECIMAL, INT, LONGING	Extract as is
DATETIME, DATETIME(6)	DATE_FORMAT(COL,'%d.%m.%Y %H:%i:%S.%f')
LONGTEXT, LONGBLOB	UPPER(MD5(COL))
VARCHAR	UPPER(MD5(COL))

O2A: 9. DATA VALIDATION

Data extraction in Oracle

<table>
<tr><td>NUMBER</td><td>replace(trim(to_char(COL,lpad('9',PRECISION,'9')||'.'||lpad('9',SCALE,'9'))),'.',decode(substr(trim(to_char(COL,lpad('9',PRECISION,'9')||'.'||lpad('9',SCALE,'9'))),1,1),'.','0.','.'))</td></tr>
<tr><td>TIMESTAMP</td><td>TO_CHAR(COL,''DD.MM.YYYY HH24:MI:SS.FF6'')</td></tr>
<tr><td>DATE</td><td>TO_CHAR(COL,'DD.MM.YYYY HH24:MI:SS')||decode(COL,null,null,'.000000')</td></tr>
<tr><td>CLOB</td><td>RAWTOHEX(DBMS_CRYPTO.Hash(to_nclob(COL),2))</td></tr>
<tr><td>BLOB, NCLOB</td><td>RAWTOHEX(DBMS_CRYPTO.Hash(COL,2))</td></tr>
<tr><td>2000+ long CHAR, VARCHAR2</td><td>RAWTOHEX(DBMS_CRYPTO.Hash(TO_CLOB(COL2))</td></tr>
<tr><td>&lt;2000 long CHAR, VARCHAR2</td><td>RAWTOHEX(DBMS_CRYPTO.Hash(UTL_RAW.CAST_TO_RAW(COL),2))</td></tr>
</table>

O2A: 9. DATA VALIDATION

```
mysql -h prod-aurora-db01-P 3306 -u pythian -p*** --batch --quick -e "SELECT
  ID,
  CUSTOMER_ACCOUNTS_ID,
  UPPER(MD5(DISCLAIMER)) DISCLAIMER,
  UPPER(MD5(TYPE)) TYPE,
  UPPER(MD5(FREQUENCY)) FREQUENCY,
  FREQUENCY_RECURRING_DAYS,
  DATE_FORMAT(CREATE_DATE,'%d.%m.%Y %H:%i:%S.%f') CREATE_DATE,
  DATE_FORMAT(LAST_UPDATE_DATE,'%d.%m.%Y %H:%i:%S.%f') LAST_UPDATE_DATE,
  UPPER(MD5(DISCLAIMER_PART2)) DISCLAIMER_PART2
FROM proddb.DATATABLE
ORDER BY 1,2,3,4,5,6,7,8,9
" | sed 's/\t/","/g;s/\n//g' | sed -e '/^$/d' -e 's/"//g' -e 's/NULL//g' -e 's/D41D8CD98F00B204E9800998ECF8427E//g'
> m_DATATABLE.csv
```

O2A: 9. DATA VALIDATION

```
sql -s pythian/***@prod-ora-db01_PROD \
<< EOF | sed 's/\t/","/g;s/\n//g' | sed -e '/^$/d' -e 's/"//g' -e 's/NULL//g' -e
's/D41D8CD98F00B204E9800998ECF8427E//g' > o_DATATABLE.csv
set feedback off echo off pagesize 0 linesize 32000 trimspool on set null NULL arraysize 5000
set sqlformat DELIMITED , " "
SELECT
  ID,
  CUSTOMER_ACCOUNTS_ID,
  case when DISCLAIMER is null then null else RAWTOHEX(DBMS_CRYPTO.Hash(TO_NCLOB(DISCLAIMER),2)) end DISCLAIMER,
  case when TYPE is null then null else RAWTOHEX(DBMS_CRYPTO.Hash(TO_CLOB(TYPE),2)) end TYPE,
  case when FREQUENCY is null then null else RAWTOHEX(DBMS_CRYPTO.Hash(TO_CLOB(FREQUENCY),2)) end FREQUENCY,
  FREQUENCY_RECURRING_DAYS,
  TO_CHAR(CREATE_DATE,'DD.MM.YYYY HH24:MI:SS.FF6') CREATE_DATE,
  TO_CHAR(LAST_UPDATE_DATE,'DD.MM.YYYY HH24:MI:SS.FF6') LAST_UPDATE_DATE,
  case when DISCLAIMER_PART2 is null then null else RAWTOHEX(DBMS_CRYPTO.Hash(TO_NCLOB(DISCLAIMER_PART2),2)) end
DISCLAIMER_PART2
FROM HMM_OWNER.DATATABLE
ORDER BY 1,2,3,4,5,6,7,8,9;
exit;
EOF
```

O2A: 9. DATA VALIDATION

Does this approach work?

Full load & offline = easy

But what about CDC and quick switchover?

Speed and Trust?

1. Validate the data 1h before the cutover
2. Record the ID's that are different
3. Record the max id from each table
4. Re-validate only the identified Ids during switchover

O2A: 10. Remove the Unnecessary Resources

Once the work is done,

Stop and delete the replication task

Remove the endpoints

(Remove the replication instance)

(Delete the source database (after the final snapshot))

We’re still running the Oracle-to-Mysql Replciation...

ISSUES WE FACED

ISSUES WE FACED

Data type Conversions via DMS

Ora: CLOB → MySQL: LONGTEXT → Ora: NCLOB

Source Data Types for Oracle

Target Data Types for MySQL

ISSUES WE FACED

Some char conversions didn’t work 1 to 1

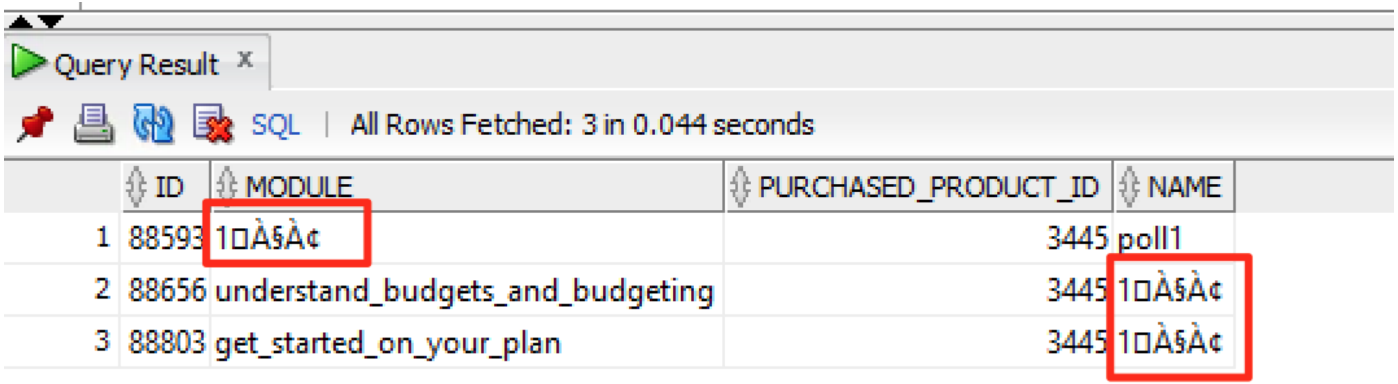

Query Result

SQL | All Rows Fetched: 3 in 0.044 seconds

	ID	MODULE	PURCHASED_PRODUCT_ID	NAME
1	88593	1□À§À¢	3445	poll1
2	88656	understand_budgets_and_budgeting	3445	1□À§À¢
3	88803	get_started_on_your_plan	3445	1□À§À¢

HEX representations:

Oracle: **3100**C380C2A7C380C2A2

Aurora: **3120**C380C2A7C380C2A2

ISSUES WE FACED

Key Collision due to Collate settings

(utf8_general_ci)

«XXX_XXcarseki@xxxxx» == «XXX_XXcarşeki@xxxxx»

ISSUES WE FACED

Aurora MySQL failed over to Read Replica

The DMS task crashed, and did not restart

We couldn’t use the cNAME, remember?

Set up the monitoring on the Replication Task Status.

ISSUES WE FACED

DMS task going into a restart loop

We’re monitoring the status changes

Binlog on MySQL side was missing

Alert was not raised

"RecoverableErrorCount" = 24 (from -1)

Use AWS CLI and json to change it

ISSUES WE FACED

11 missing rows in a multimilion-row table

?

ISSUES WE FACED

Downsized jpg stored in a BLOB column in MySQL

?

SUMMARY

SUMMARY

- DMS works
- Price/performance is awesome
- Become familiarized with the AWS CLI
- Think about Monitoring (CW)
- Validation trouble (at the time of migration)
- Visibly a very important product of AWS / features added often
- SCT simplifies schema creation

THANK YOU

?

@MarisDBA

elsins@pythian.com