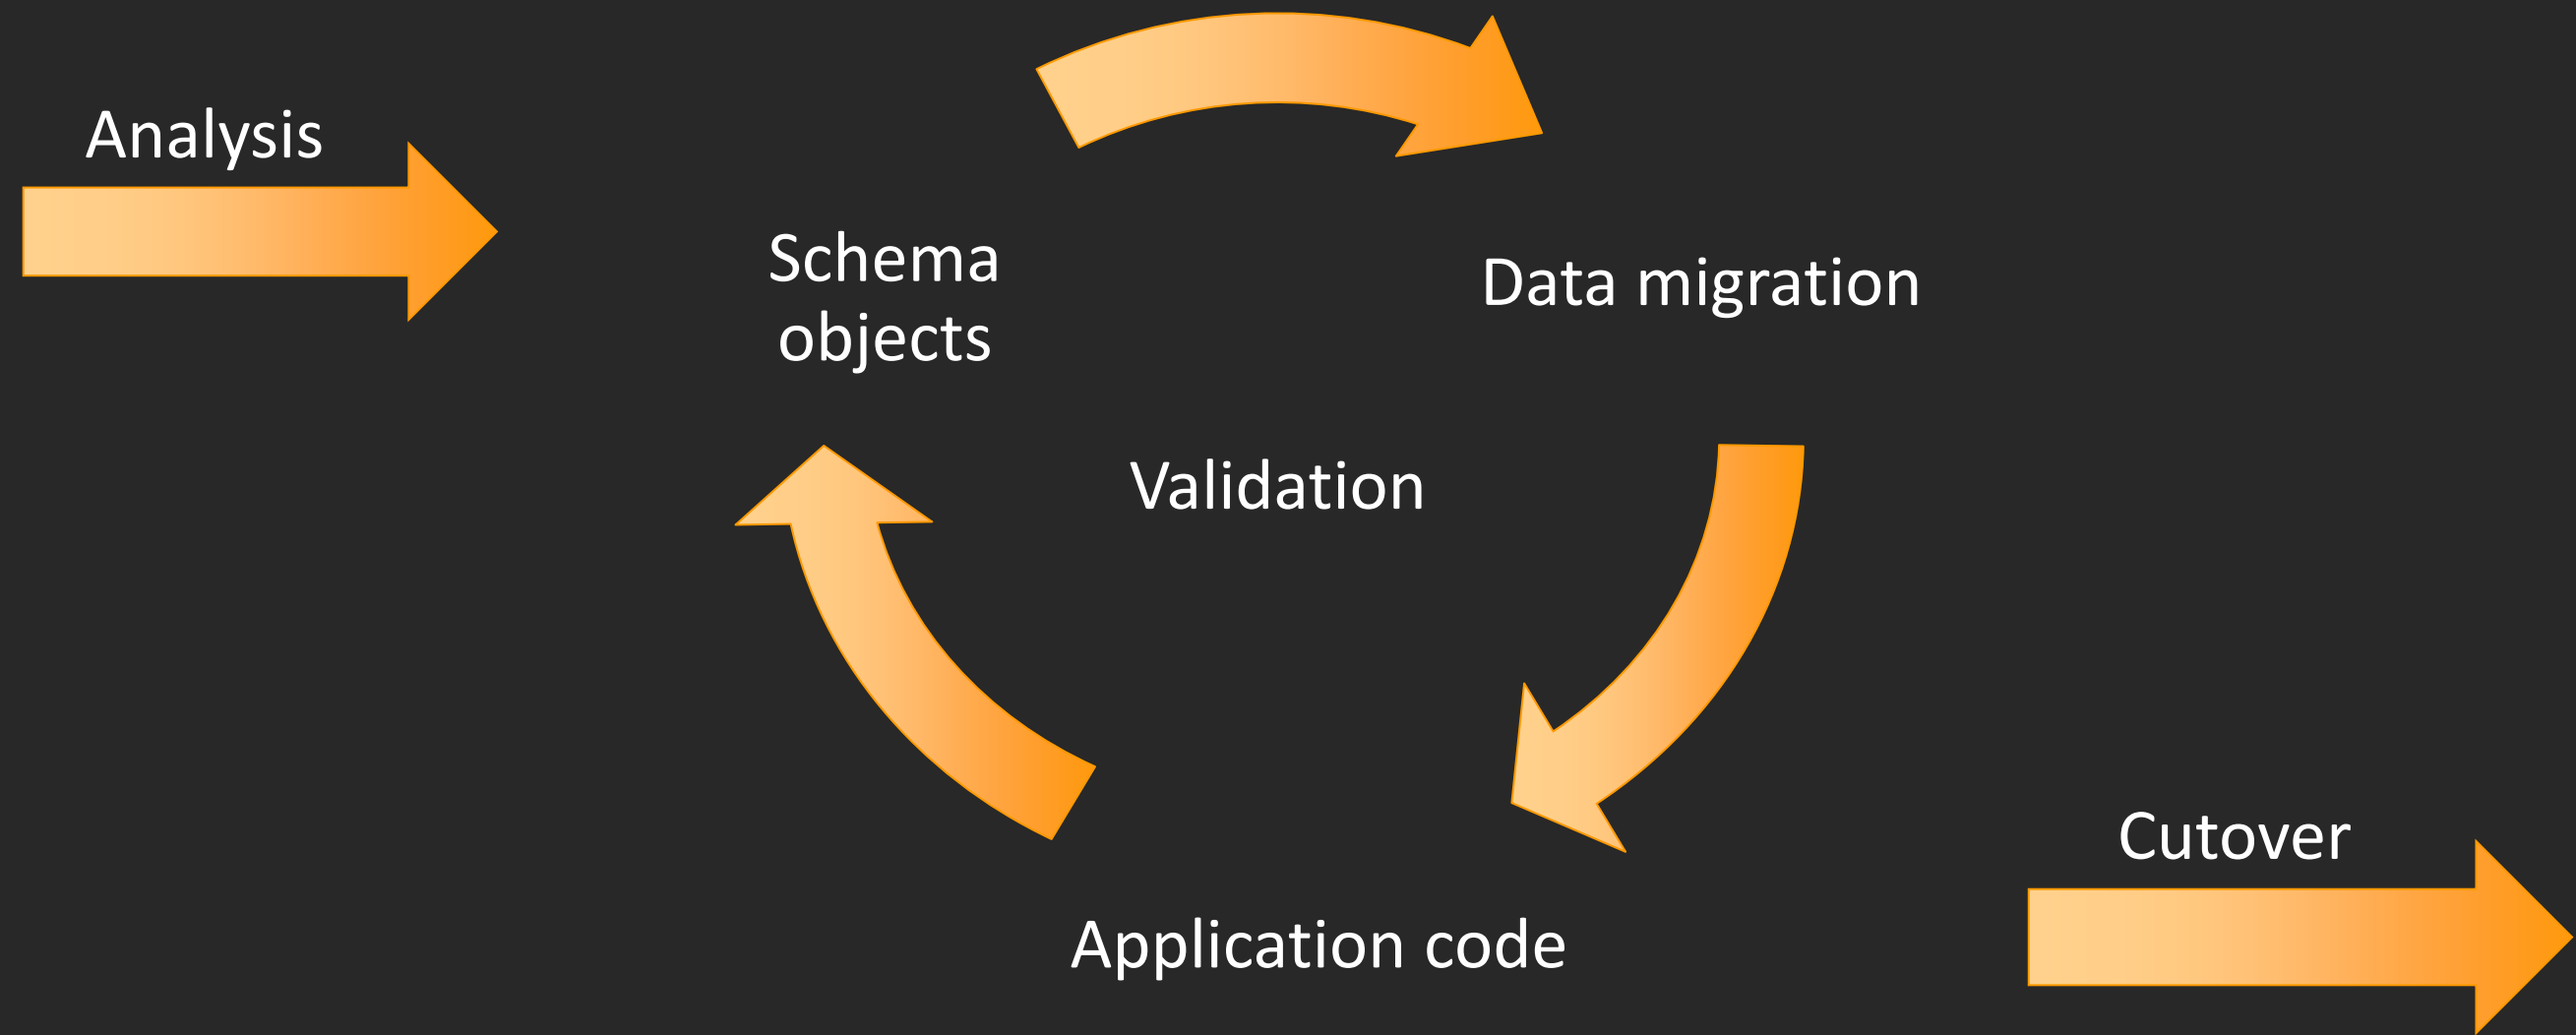




AWS 데이터베이스 서비스로 마이그레이션하기

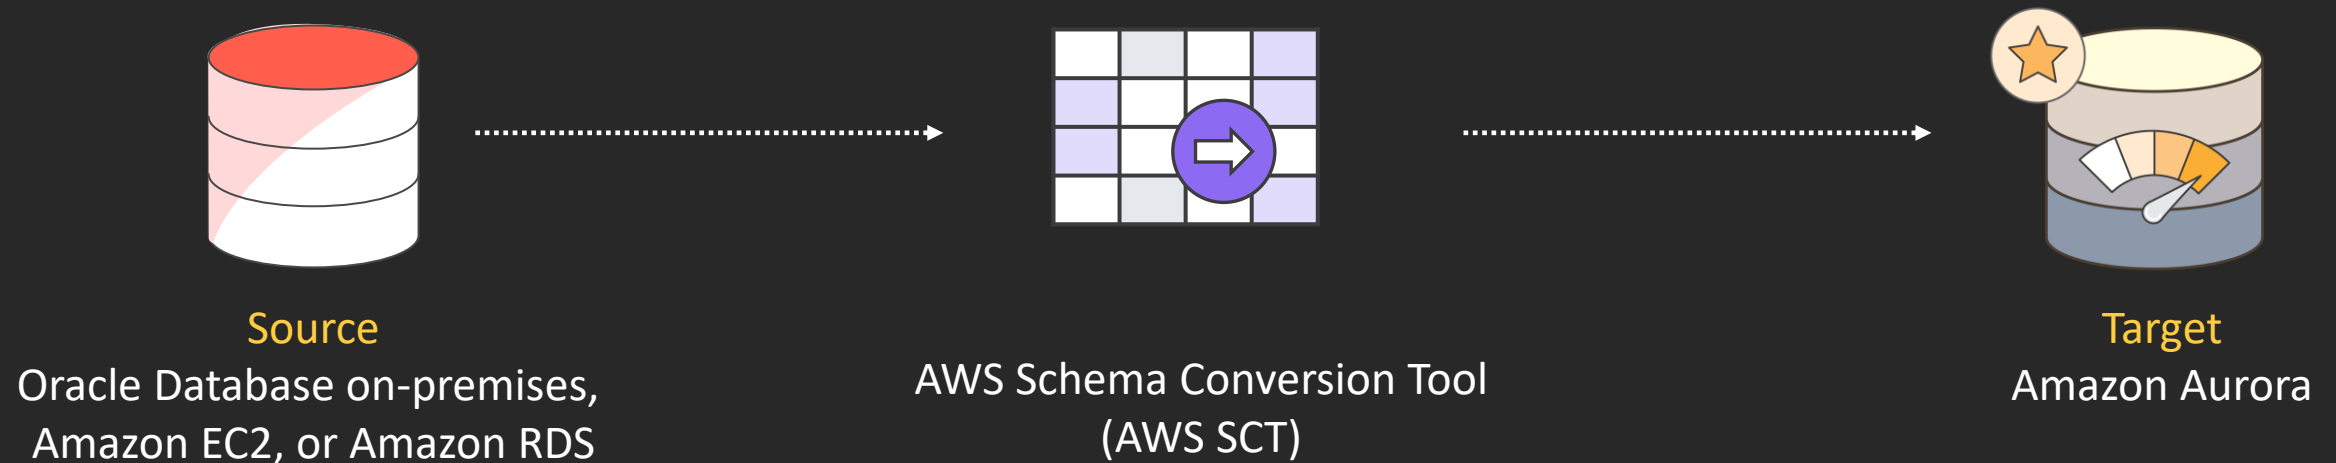
Jihoon Kim, Database Specialist Solution Architect
Dec, 2021

A typical database migration lifecycle

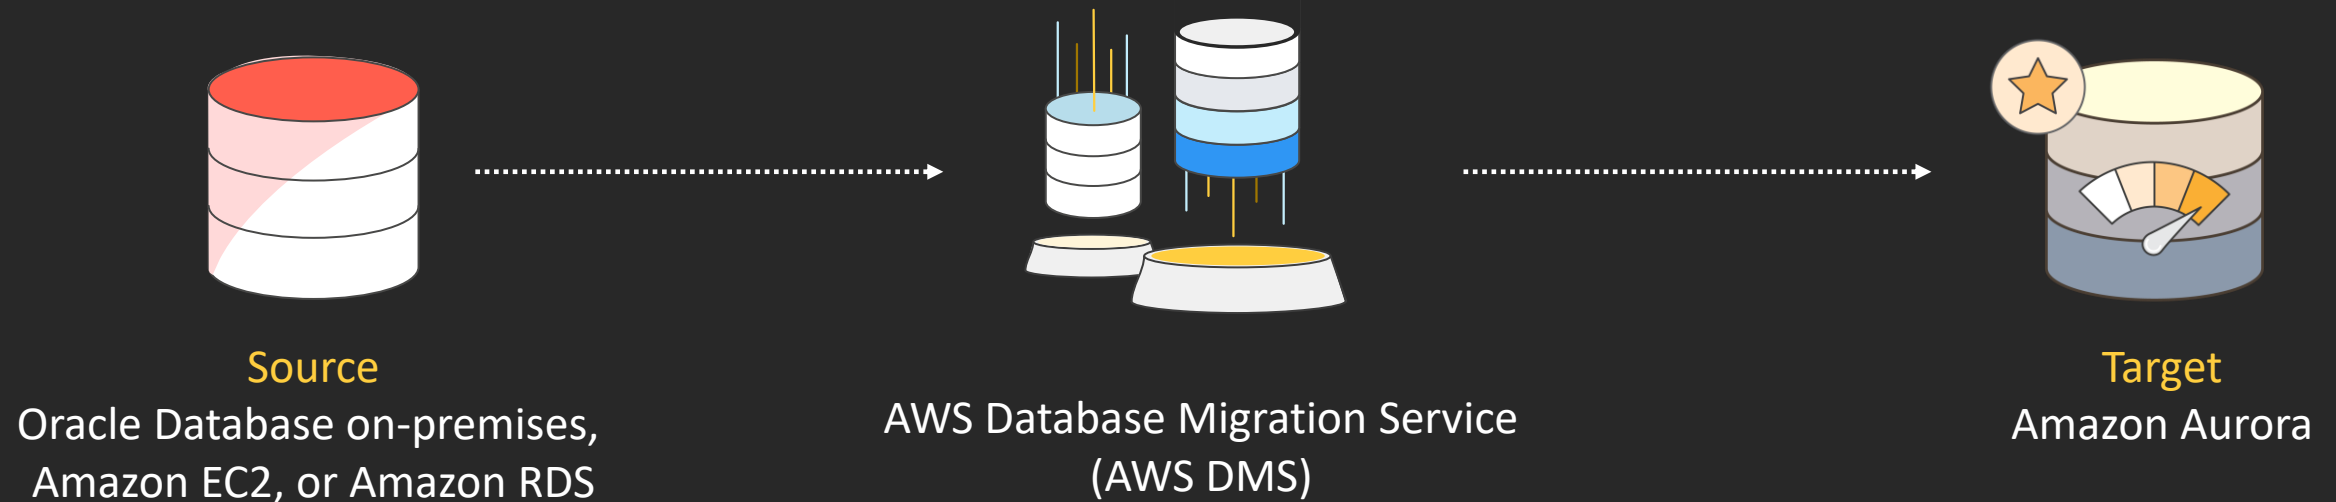


Database migration process

Step 1



Step 2



Oracle to Amazon Aurora migration playbook

- In-depth guidance on Oracle to Aurora PostgreSQL migration
- With detailed examples
- Migration best practices

Schema



AWS SCT

Data



AWS DMS

Best practices



Playbook

	Oracle Feature	PostgreSQL Feature	Compatibility
Link	Index Organized Tables (IOTs)	PostgreSQL "Cluster" Tables	Yes*
Link	Common Data Types	Common Data Types	Yes
Link	Table Constraints	Table Constraints	Yes
Link	Table Partitioning including: RANGE, LIST, HASH, COMPOSITE, Automatic LIST	Table Partitioning including: RANGE, LIST	Yes*
Link	Exchange & Split Partitions	N/A	None
Link	Temporary Tables	Temporary Tables	Yes*
Link	Unused Columns	ALTER TABLE DROP COLUMN	Yes
Link	Virtual Columns	Views and/or Function as a Column	Yes*
Link	User Defined Types (UDTs)	User Defined Types (UDTs)	Yes
Link	Read Only Tables & Table Partitions	Read Only Roles and/or Triggers	Yes*
Link	Index Type	Link Recovery Manager (RMAN)	Link AWS Aurora Snapshots
Link	B-Tree Index	Link Flashback Database	Link AWS Aurora Snapshots
Link	Composite Index	Link 12c Multi-tenant architecture: PDBs and CDB	Link Databases
Link	BITMAP Index	Link Tablespace & DataFiles	Link Tablespaces
Link	Function-based Index	Link Data Pump	Link pg_dump & pg_restore
Link	Global and Local Indexes	Link Resource Manager	Link Separate AWS Aurora Clusters
Link	Identity Column	Link Database Users	Link Database Roles
Link	MVCC (Table & File)	Link Database Roles	Link Database Roles
Link	Character Set	Link SGA & PGA Memory	Link Memory Buffers
Link	Transaction	Link V\$ Views & Data Dictionary	Link System Catalog Tables, Statistics Collector, AWS Aurora Performance Insights
Link		Link Log Miner	Link Logging Options
Link		Link Instance & Database Parameters (SPFILE)	Link AWS Aurora Parameter Groups
Link		Link Session Parameters	Link Session Parameters
Link		Link Alert.log (error log)	Link Error Log via AWS Console
Link		Link Automatic and Manual Statistics Collection	Link Automatic and Manual Statistics Collection
Link		Link Viewing Execution Plans	Link Viewing Execution Plans

<https://d1.awsstatic.com/whitepapers/Migration/oracle-database-amazon-aurora-postgresql-migration-playbook-12.4.pdf>

AWS Schema Conversion Tool (AWS SCT)

The AWS Schema Conversion Tool helps automate many database schema and code conversion tasks when migrating from source to target database engines

- Convert database schema
- Convert data warehouse schema
- Convert application SQL



Components of the AWS SCT console

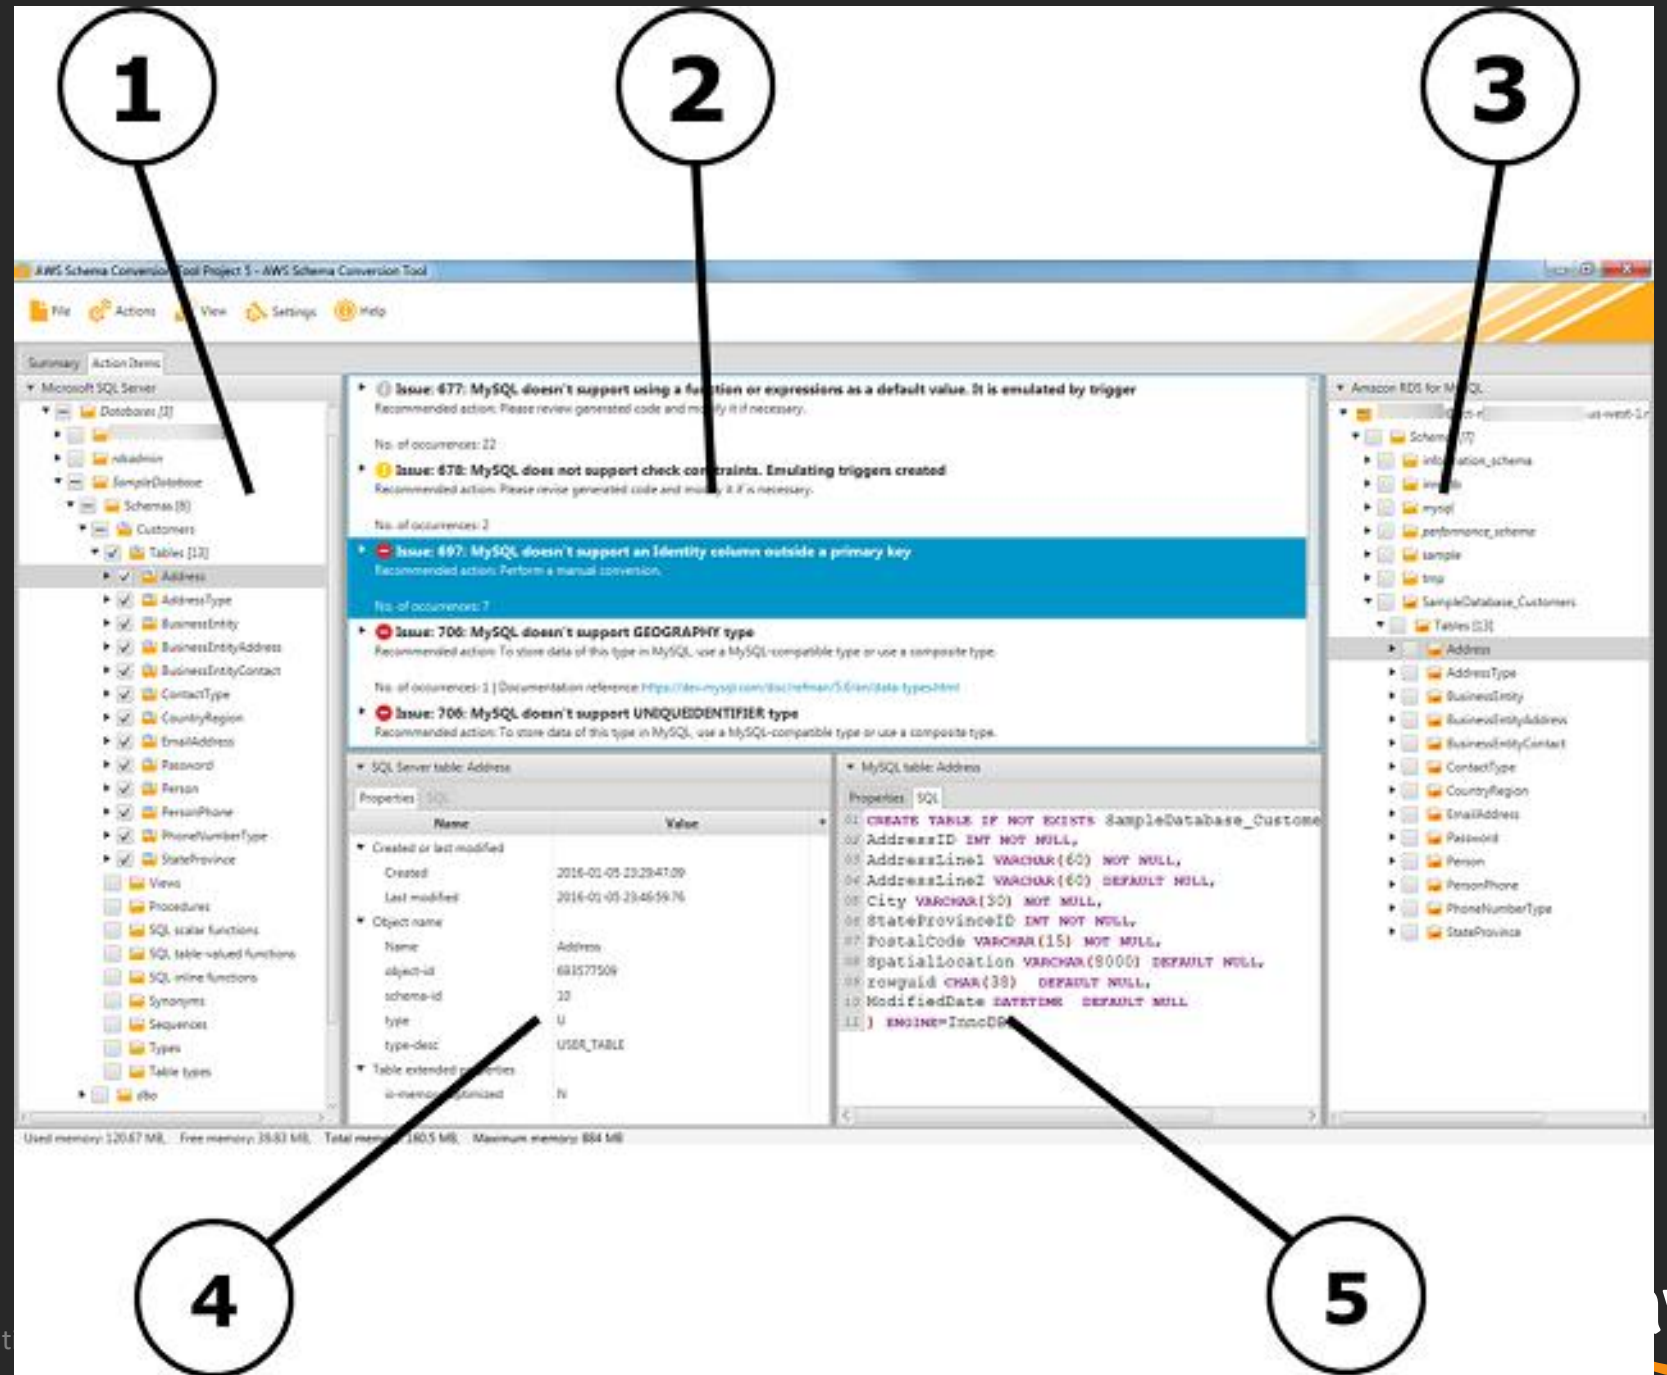
1. Source schema

2. Action items

3. Target schema

4. Schema element details

5. Edit window



AWS SCT can tell you how hard the migration will be

Database Migration Assessment Report

Source Database: C:\oracledb\product\11.2.0.3.0\sqlplus\sqlplus.exe
Oracle Database 11g 11.2.0.3.0 (64bit Production), Standard edition



Executive Summary


We completed the analysis of your Oracle source database and estimate that 76% of the database storage objects and 65% of database code objects can be converted automatically or with minimal changes if you select Amazon Aurora (PostgreSQL compatible) as your migration target. Database storage objects include schemas, tables, table constraints, indexes, types, sequences, synonyms, view constraints and clusters. Database code objects include triggers, views, materialized views, materialized view logs, procedures, functions, packages, package constants, package cursors, package exceptions, package variables, package functions, package procedures, package types and package collection types. Based on our analysis of SQL syntax elements of your source database schema, we estimate that 98% of your entire database schema can be converted to Amazon Aurora (PostgreSQL compatible) automatically. To complete the migration, we recommend 606 conversion action(s) ranging from simple tasks to medium-complexity actions to significant conversion actions.

Database Objects with Conversion Actions for Amazon Aurora (PostgreSQL compatible)

Of the total 327 database storage object(s) and 57 database code object(s) in the source database, we were able to identify 250 (76%) database storage object(s) and 37 (65%) database code object(s) that can be converted to Amazon Aurora (PostgreSQL compatible) automatically or with minimal changes.

77 (24%) database storage object(s) require 77 significant user action(s) to complete the conversion.

20 (35%) database code object(s) require 20 significant user action(s) to complete the conversion.



Oracle Database 11g 11.2.0.3.0 (64bit Production), Standard edition

Migrations

If you migrate your Oracle database to Amazon Aurora (PostgreSQL compatible), we recommend the following actions.

Storage Object Actions

Table Changes

Not all tables can be converted automatically. You'll need to address these issues manually.

Issue 5213: PostgreSQL expands fractional seconds support for TIME, DATETIME, and TIMESTAMP values, with up to microseconds (6 digits) of precision

Recommended Action: Review your transformed code and modify it if necessary to avoid a loss of accuracy.

Issue Code: 5213 | Number of Occurrences: 223 | Estimated Complexity: Simple

Documentation References: <http://www.postgresql.org/docs/9.6/static/datatype.html>

```
Schemas.CPS_DICOM.Tables.CPS_ACTION.Columns.SYS_TS
Schemas.CPS_DICOM.Tables.CPS_ACTION_H.Columns.ORIG_SYS_TS
Schemas.CPS_DICOM.Tables.CPS_ACTION_H.Columns.SYS_TS
Schemas.CPS_DICOM.Tables.CPS_ACTION_PARAMETER.Columns.SYS_TS
Schemas.CPS_DICOM.Tables.CPS_ACTION_PARAMETER_H.Columns.ORIG_SYS_TS
+218 more
```

Issue 5340: Unable to convert functions

Recommended Action: Use suitable function or create user defined function.

Issue Code: 5340 | Number of Occurrences: 77 | Estimated Complexity: Significant

Schemas.CPS_DICOM.Tables.CPS_ACTION.Columns.SYS_TS: 0:28
Schemas.CPS_DICOM.Tables.CPS_ACTION_H.Columns.SYS_TS: 0:28
Schemas.CPS_DICOM.Tables.CPS_ACTION_PARAMETER.Columns.SYS_TS: 0:28
Schemas.CPS_DICOM.Tables.CPS_ACTION_PARAMETER_H.Columns.SYS_TS: 0:28
Schemas.CPS_DICOM.Tables.CPS_ADMITTING_DIAGNOSIS.Columns.SYS_TS: 0:28
+72 more

! **Issue 5554: PostgreSQL doesn't support virtual columns**

Recommended Action: They were emulated by a trigger.

Issue Code: 5554 | Number of Occurrences: 7 | Estimated Complexity: Simple

Documentation References: <http://www.postgresql.org/docs/9.6/static/sql-createtable.html>

Schemas.CPS_DICOM.Tables.CPS_ENCAPSULATED_DOC
Schemas.CPS_DICOM.Tables.CPS_EQUIPMENT
Schemas.CPS_DICOM.Tables.CPS_PERSON_IDENT
Schemas.CPS_DICOM.Tables.CPS_PERSON_NAME
Schemas.CPS_DICOM.Tables.CPS_SERIES

Connect AWS SCT to
source and target
databases

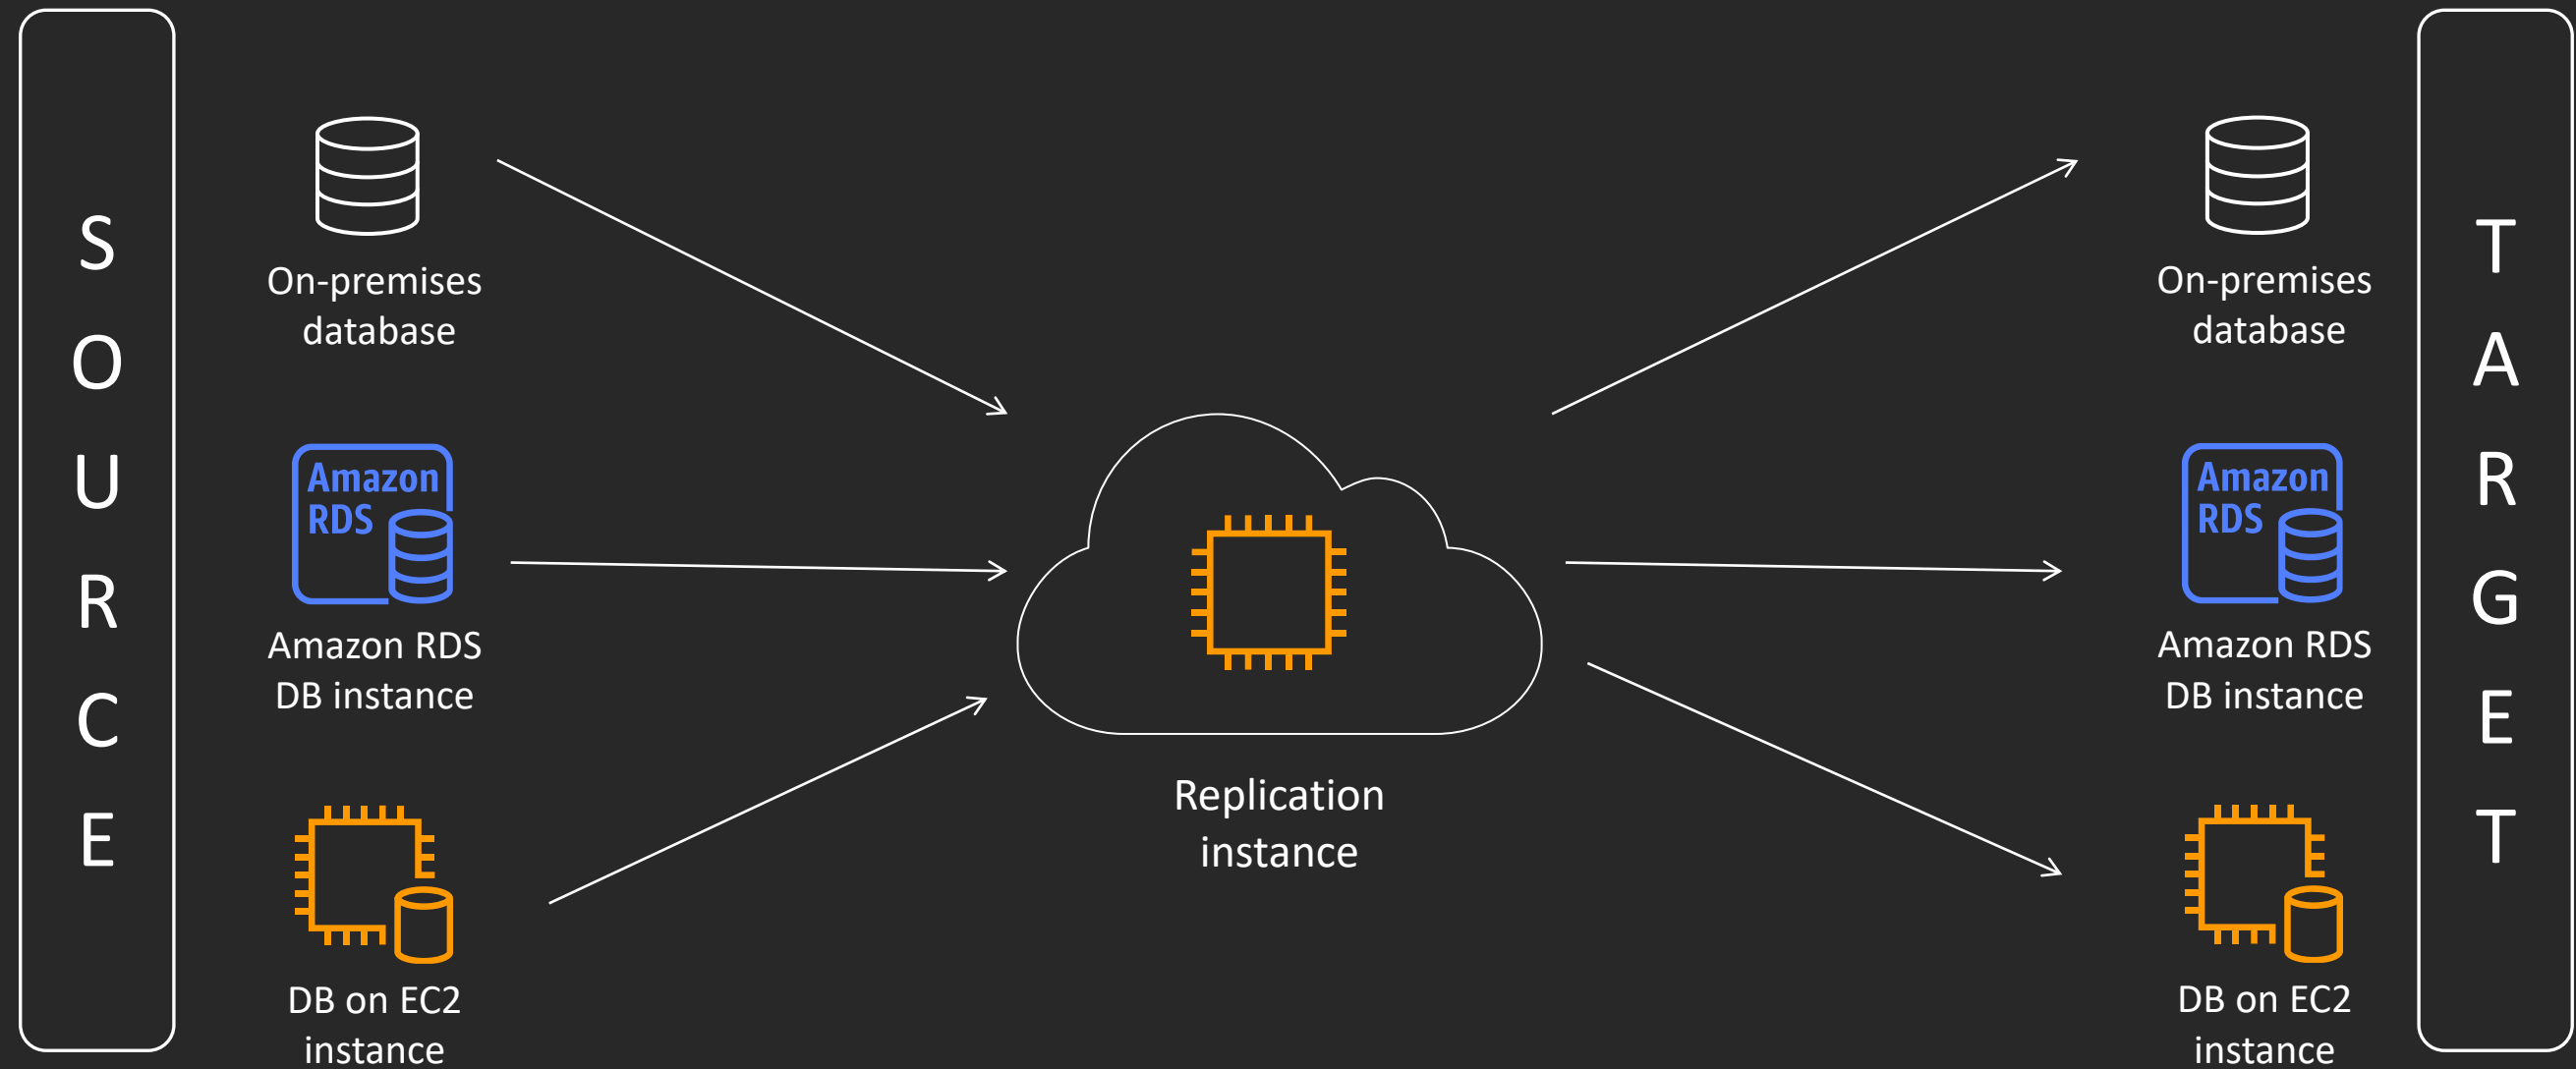
Run assessment report

Read executive summary

Follow detailed
instructions



AWS Database Migration Service (AWS DMS)



AWS DMS components—Tasks and replication instances

- **Tasks** run on a **replication instances**
- Contain two **endpoints** (source and target)
- Specify **selection and transformation rules**
- Replication instances support multiple tasks




Rules and filters


- **Selection rules—Applied at the source**
 - Include or exclude schemas and tables
 - Filter rows based on column values
 - Can use multiple filters
- **Transformation rules—Applied on the target**
 - Rename schemas, tables, and columns
 - Remove columns, change case
 - Add/remove prefix or suffix



Migrate existing data (full load)



Change data capture (CDC)



Database migration procedure

1. **Convert the schema** using AWS SCT
2. **Drop foreign keys and indexes** on the target database
3. Set up a **AWS DMS** task to migrate initial data (full load)
4. **Re-create foreign keys and indexes** on the target database after the full load task is complete
5. Set up another **AWS DMS** task to **capture and apply data changes (CDC)** from the source
6. **Point applications** to the target database

Migration best practices

Planning your migration

- Migration **complexity analysis**
 - Schema complexity analysis
 - Redo rate and TPS on source
 - Analyze the impact due to feature differences
 - Be aware of common migration issues
(ex. https://wiki.postgresql.org/wiki/Oracle_to_Postgres_Conversion)
- Identify the **pilot/quick win candidates**

Schema migration complexity analysis

Factor	Weightage	Comment
Tables	0.5	Non-partitioned tables will be converted automatically. Look for data type, character set issues.
Materialized views	3	Only fully refreshed Mviews are supported
Triggers	3	Review or rewrite the logic
PL/SQL procedures	4	Review or rewrite the logic
PL/SQL packages	15	Packages with ≥ 3 procedures or functions, review, or rewrite the logic
DB_link	1	Convert DBLinks using postgres_fdw extension
Partitions	0.1	Create all the child partitions manually. Apply the partitioning architecture with the help of functions/triggers.
LOB	1	LOB columns requires special handling

Schema migration complexity analysis—Example

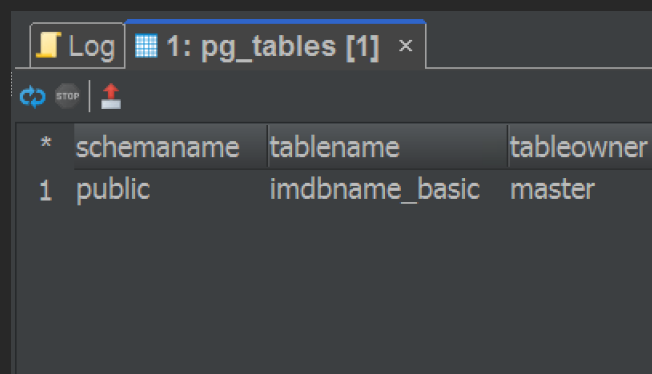
DB_NAME	TAB	MV	DLINK	PRC	FUNC	PKG	TRG	Partitions	LOB	MAX_REDO_HR	TOT_WEIGHTAGE
DB1	'40:20'	'0:0'	'3:3'	'37:148'	'0:0'	'5:75'	'14:42'	'50:5'	'0:0'	0.76	327

Factors	SMALL	MEDIUM	LARGE	XLARGE
Complexity	1–500	500–1000	1000–2000	2000+
Redo_GB_Per_Hour	0–4	4–10	10–20	20+

Number of DBs	SMALL	MEDIUM	LARGE	XLARGE
Org1	20	10	5	2
Org2	30	5	5	5



Understand basic database engine differences



A screenshot of the PostgreSQL pg_tables view. The table has three columns: schemaname, tablename, and tableowner. The first row shows 'public' as the schema, 'imdbname_basic' as the table name, and 'master' as the owner.

	schemaname	tablename	tableowner
1	public	imdbname_basic	master


PostgreSQL is a lowercase data dictionary

```
1. EXCEPTIONS = INSERT E...  
    RETURN - 1;  
    END IF;  
    END;  
    WHEN others THEN  
    RETURN - 1;  
;  
DY$  
GUAGE plpgsql;
```

Use “exception handlers” when needed, not by default

- B-Tree
- Generalized Inverted Index (GIN)
- Generalized Inverted Search Tree (GiST)
- Space partitioned GiST (SP-GiST)
- Block Range Indexes (BRIN)
- Hash

PostgreSQL has six index types



Store your BLOBs in Amazon S3 instead of the database

Examples

Set the schema search path:

```
SET search_path TO my_sc
```


search_path replaces PUBLIC SYNONYM

ORACLE DATATYPES

VARCHAR2	CHAR
NUMBER	DATE
TIMESTAMP	BLOB
CLOB	ROWID

PostgreSQL has 64 datatypes

Keep your apps running during migration



Start replication instance

Connect to source and destination

Select tables, schemas, or databases




Application users

Let AWS DMS create tables, load data, and keep them in sync

Switch applications to new target at your convenience


AWS DMS performance factors

- Resource availability on the source
- Network throughput
- Replication server capacity
- Number of objects, large tables
- LOB controls and performance



Large objects (LOBs) performance

- **LOB controls**
 - “Don’t include”—omit LOB columns
 - “Limited LOB mode”—specify “max LOB size”
 - “Full LOB mode”—specify “LOB chunk size”
- **LOB performance**
 - Whenever possible use “limited LOB mode”



Performance with high TPS source systems

Transactional apply (default)

Maintains transaction commit order


Batch apply

Build a batch with changes from source

Intermediate net changes on target


Condense changes and apply

- Applies all inserts, all updates, and all deletes in the same order

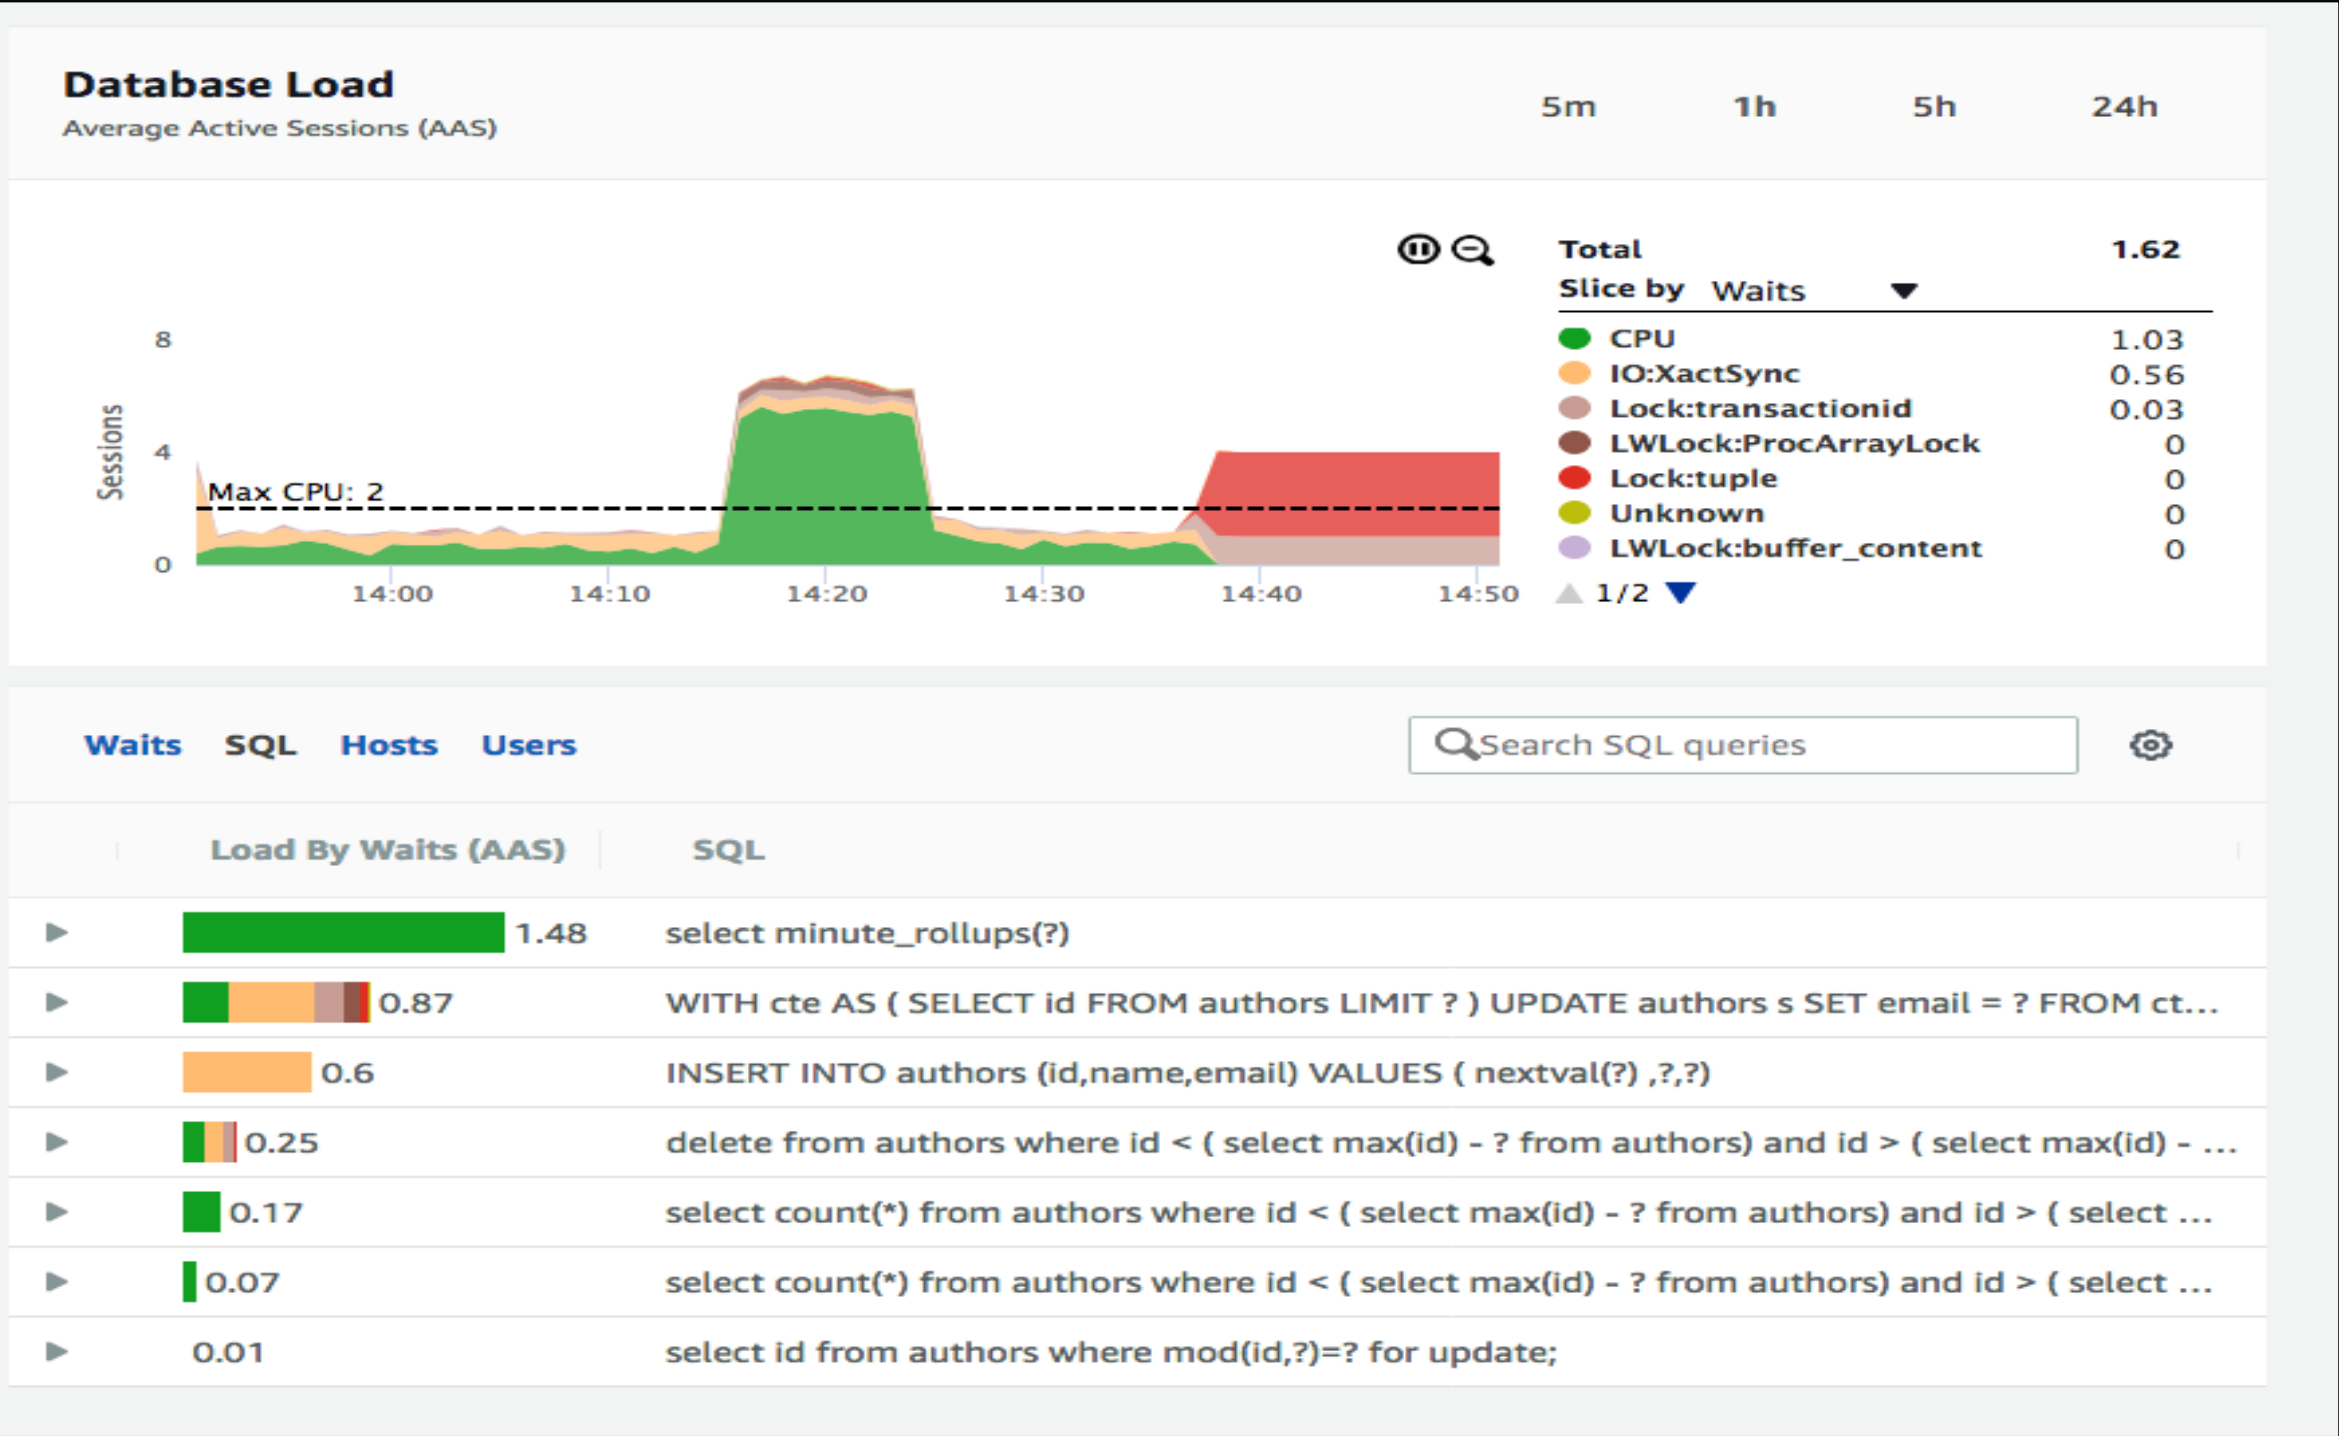


Performance testing

- Load testing
- Identify critical application features
- Features with known performance issues
- Identify Top SQLs/complex SQLs



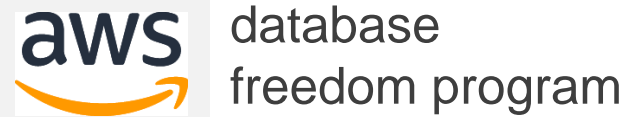
Performance insights



How AWS can help



Tools



Programs



Partners

AWS database migration partners



Thank You